

Real Analysis

by Anshula Gandhi



About this book

"What is the point of this?" is a common question in math classes for a reason.

We often learn math concepts before understanding the philosophical questions that motivated them. For example, I learned about bijections before I realized it answered the question: *are some infinities bigger than others?* All of the abstract math textbooks I could find prioritized comprehensiveness and exposition over why the problems were fun and interesting in the first place. Luckily, I found a few friends and teachers who taught me exactly why these subjects were so fascinating, and knew the interesting questions to ask, and once I knew that, I could turn to my textbooks for comprehensiveness and exposition.

This comic is meant to teach real analysis, a subject that often serves as an introduction to abstract math. It's the textbook I wish I had - one that includes visuals, stories, and the philosophical questions that lead to core mathematical concepts. After all, math is like philosophy, but with answers.

I hope this book is something you'd want to read whether you need to pass an analysis class, or pass some time. Either way, I hope you enjoy reading about the subject that brought so much joy and comfort to me.

Real Analysis

Chapter 0
Introduction

What is real analysis?

You learned some amount of rubbish about calculus in high school (even if you didn't realize it).



Analysis is about cleaning up that rubbish.




To do analysis...




We have to forget everything we know about math.



Then we'll build it back up again, proving each step, making sure it's all true.



Note: Analysis is about *proving* calculus, not *doing* calculus.



For example, you'll *prove* that integrals exist, but never *calculate* a single one.



Why do analysis?

Disclaimer: Analysis isn't technically "useful," in that:
1) You can rarely apply it to the real world, and
2) It probably won't help you land a job.



But despite analysis being close to useless in a job application, it is useful in that:

1) The hand-waving of calculus can lead to paradoxes. Analysis fixes that by proving calculus rigorously.



2) Analysis presents an opportunity to think and prove things in an entirely new way.



3) Analysis offers escape from reality - a chance to philosophize about problems that have nothing to do with your everyday life.



And sure, analysis does have some practical applications. But who cares, anyway? We do it because it's fun.



DIGRESSION

Why do analysis books rarely use pictures?

Analysis came about as part of a movement in the 1700s. Some mathematicians wanted math to be more abstract and "pure," and therefore to rely less on the "crutches" of figures and diagrams.

It was something of a challenge, perhaps, to define math without relying on figures.

And that's probably why most analysis texts shy away from visuals: because a big point of analysis was to not need visuals anymore.

Real Analysis

Chapter 1
Number Systems

How do you prove that something doesn't exist?

Proving something *doesn't* exist can be a lot harder than proving something *does* exist.

A sociologist, Jerry Lembcke, ran into that problem.



He'd heard stories about antiwar protestors spitting on Vietnam War soldiers as they returned home to America.



But as Lembcke looked into the phenomenon, he could find no single instance of this spitting.



But he couldn't say for sure that "the spitting never happened."



If he only had to prove the spitting happened, he would just have to find one account of it.



But how could he prove that nobody ever spat on a war veteran?



It would be impossible to interview every single Vietnam War veteran - dead and alive.



The sociologist admitted that he couldn't say the "spitting protestor" phenomenon was untrue. He could only say he found no evidence it was true.

How do you prove non-existence in math?

So, people say "you can't prove a negative statement." Not true. It's hard. But in math, it's possible with a "proof by contradiction."

We start by assuming something *does* exist, and reason through it.



If we find a contradiction within our reasoning, we must conclude that thing *does not* exist.




But why?

When two people contradict each other, you know one of them has got to be wrong.



Similarly, in math, if two of the lines in your proof contradict each other, then there's a lie in there somewhere.



And if every step you took in the proof was correct, then the lie can't be in any of the steps you took. The lie must be way back in the very first assumption you made.

How do you prove irrationality?

When an ancient Greek discovered that some numbers are irrational*...

He was exiled.
(It was heresy to say numbers were so disorderly.)




But how could he have proven that irrational numbers exist?



*Irrational numbers are those that can't be written as a ratio of integers (e.g. pi).

Did he line up all the infinite fractions in the world, and compare each to a number he thought was irrational?



He couldn't have. He'd have to have checked all infinite fractions in the world before ensuring none of them was his guy.



That method of proof would be impossible.

Instead, we could use proof by contradiction.




Is the square root of two irrational?

First, let's declare that $\sqrt{2}$ exists.
(We should be explicit about the rules we're going by.)

Let's assume, for the purpose of contradiction, that $\sqrt{2}$ is rational. (It's not.)



Then there must exist a fraction $\frac{p}{q}$ that equals $\sqrt{2}$.




Let's use a p and q that have no common factors (so that $\frac{p}{q}$ is simplified.)




Now we'll show the contradiction: even though we chose p and q to share no common factors, they always end up sharing a factor of two.

Let's interrogate p and q separately.

Let's prove p is divisible by two.



Let's prove q is divisible by two.



So, p and q share
a factor of two.
Contradiction!



* Remember, we chose a p
and q that shared no
common factors.

But all
the
steps we
took
were
correct...

So the flaw must be in
our assumption, and the
only thing we assumed
was that $\sqrt{2}$ is rational.

So, $\sqrt{2}$ can't be rational.




But do irrational numbers actually exist?

A friend once told me that irrational numbers don't exist in the real world.

"What, no," I said. "Of course they do."




"If I draw a triangle with sides of length 1..."



"Then the hypotenuse has length $\sqrt{2}$. That's irrational."



"Well, no," she replied. "If you draw that length in the real world, that line stops somewhere. If you zoom in close enough and use a ruler, you'll see that the decimal digits of that length don't go on forever."






"Even that side of length 1 you drew isn't exactly length 1. The probability that length is exactly 1.00000000, as you add on infinite 0s after the decimal point, diminishes to 0 percent."






So why bother with all of this stuff if it all only exists in our heads?



How does proving irrationality relate to analysis?

So what does proving that irrational numbers exist have to do with real analysis - the building of calculus?

The purpose of proving that irrational numbers exist is to show that there are 'gaps' in the rational line of numbers.




This gap is somewhat surprising, since it seems that rationals are densely packed. That is, between every two rational numbers, you can find another rational number (consider the number $(p+q)/2$ that exists between rationals p and q). So, given that rationals are so dense, it's surprising that we found a gap at the square root of two.



It's not only surprising, but also somewhat inconvenient that rationals have gaps.




Sets that don't have gaps (or 'complete' sets), such as the real line, are useful for building up calculus.*




*For example, we know that limits are a foundational concept in calculus. But a sequence might not have a limit in an incomplete set. For example, consider the sequence of rational numbers that slowly approaches π : 3, 3.1, 3.14, and so on. It will have no limit in the rationals (because its limit is π).

Why does analysis feel so unrelated to calculus?

Analysis is presented from the bottom up - starting from basic axioms and building towards calculus.



But historically, analysis was developed from the top down - starting from calculus and going down towards the basic axioms.



So this is why you'll start analysis and learn about things like 'irrationality' and 'gaps' and be thinking: what the heck does this have to do with calculus? It becomes clear only later.



And so, next we'll next discuss ordered sets.



Yeah, it's not immediately clear how ordered sets help us build up calculus.

And yeah, it is totally different than what we just discussed - irrationality.

But let's just embrace the chaos.



How do we know where to put numbers on a number line?

It's useful to know that some numbers are bigger than others.



But as of now in our building-up of calculus, we haven't yet introduced any concept of some numbers being bigger than others.

We don't have big numbers and small numbers, right now.



We just have numbers.



We say that 5 is 'bigger' than 3 only because people from a long time ago said so. That's how they decided to order 5 and 3.

And so, to know which numbers are bigger... we need to define the order they go in.



Let's say that a 'proper' way to order numbers is any way such that if you compare two of them, one is equal to, bigger than, or lesser than the other. More precisely, when comparing numbers 'a' and 'b' in an ordered set, exactly one of the following must hold:



*For example, a 'proper' ordering for the set $\{3, 5\}$ can be that 3 equals 5 (it's not the ordering that most mathematicians go by, but it does satisfy our definition of a proper ordering).

But an order can't be that 3 is both greater than and less than 5. And it can't be that 3 has no relation to 4. If these last two were orderings, it would be impossible to answer the question: which is bigger: 3 or 4?

First, what exactly do I mean by *order*? I mean that given two numbers, the *order* should tell us *which number is bigger*.
For example, we can ask...

How should we order the set $\{1, 2, 3\}$?

We could order it like:
 $1 < 2 < 3$.



Or, we could order it like:
 $1 > 2 > 3$.




(This is a 'proper' ordering, according to our definition. It is just not the one that most mathematicians use.)

Now for a tougher ordering problem.
How should we order the rationals?

The ordering can't just say that the bigger number is bigger.

The purpose of ordering is to *define* the word 'bigger'!


Instead, let's say our ordering is that the rational s is bigger than r when $s-r$ is positive*.



*Sounds good right? But we haven't defined positive. Now, in your head you might be thinking that it's *obvious* how to define positive: a number is positive when it's bigger than zero. But then, our definition of *positive* relies on *bigger* (*positive* means a number is *bigger* than zero) and our definition of *bigger* relies on *positive* (a number is *bigger* than another when $s-r$ is *positive*). So for now, I'll trust we all know the difference between a positive and negative number. But we will formally define what positive and negative mean later.

Ok. Now we know how to order sets. So, we can figure out which numbers are bigger than others.

Does every set have a biggest number?



Consider the set containing numbers (let's say they're all positive integers) representing how cool every person's mom is.



The bigger the number, the cooler the mom.

That set does have a biggest number - the coolness of my mom.



But do you think every set of items has a biggest number?

Sure, sets with finite elements (like our mom-coolness set) contains a biggest number.



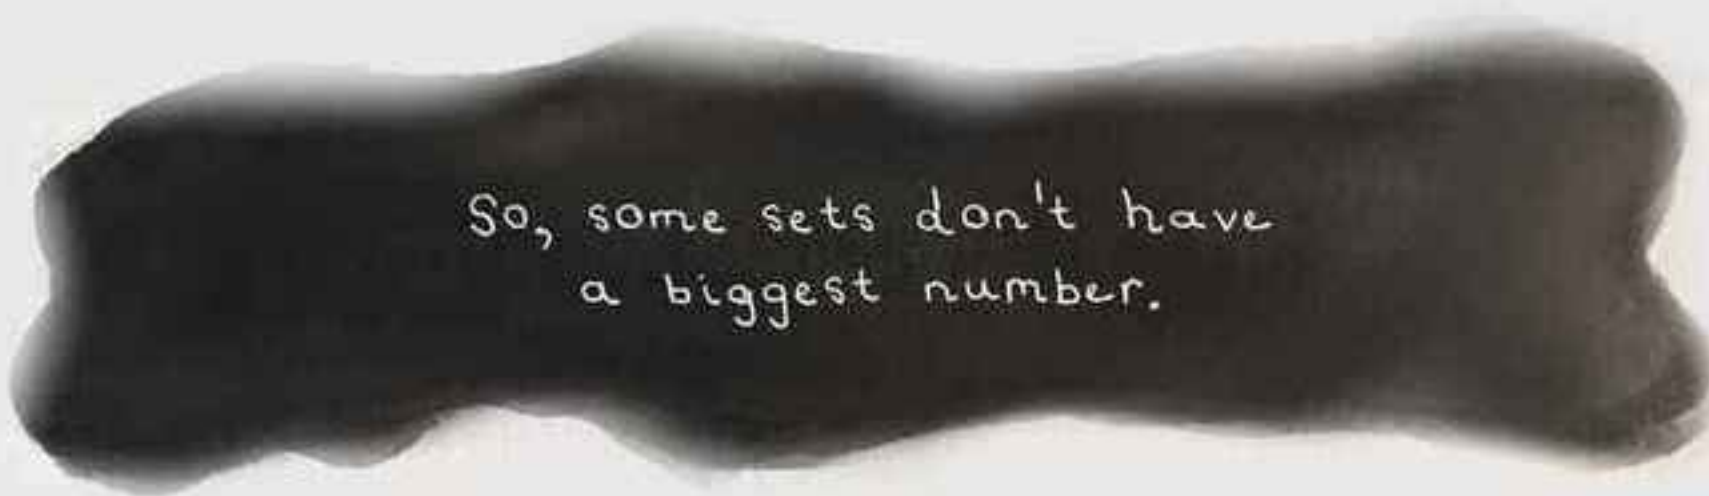
But imagine a world with infinite moms, each mom cooler than the last. (That is, their coolnesses make up the set $\{1, 2, 3, 4, 5, \dots\}$ and so on).



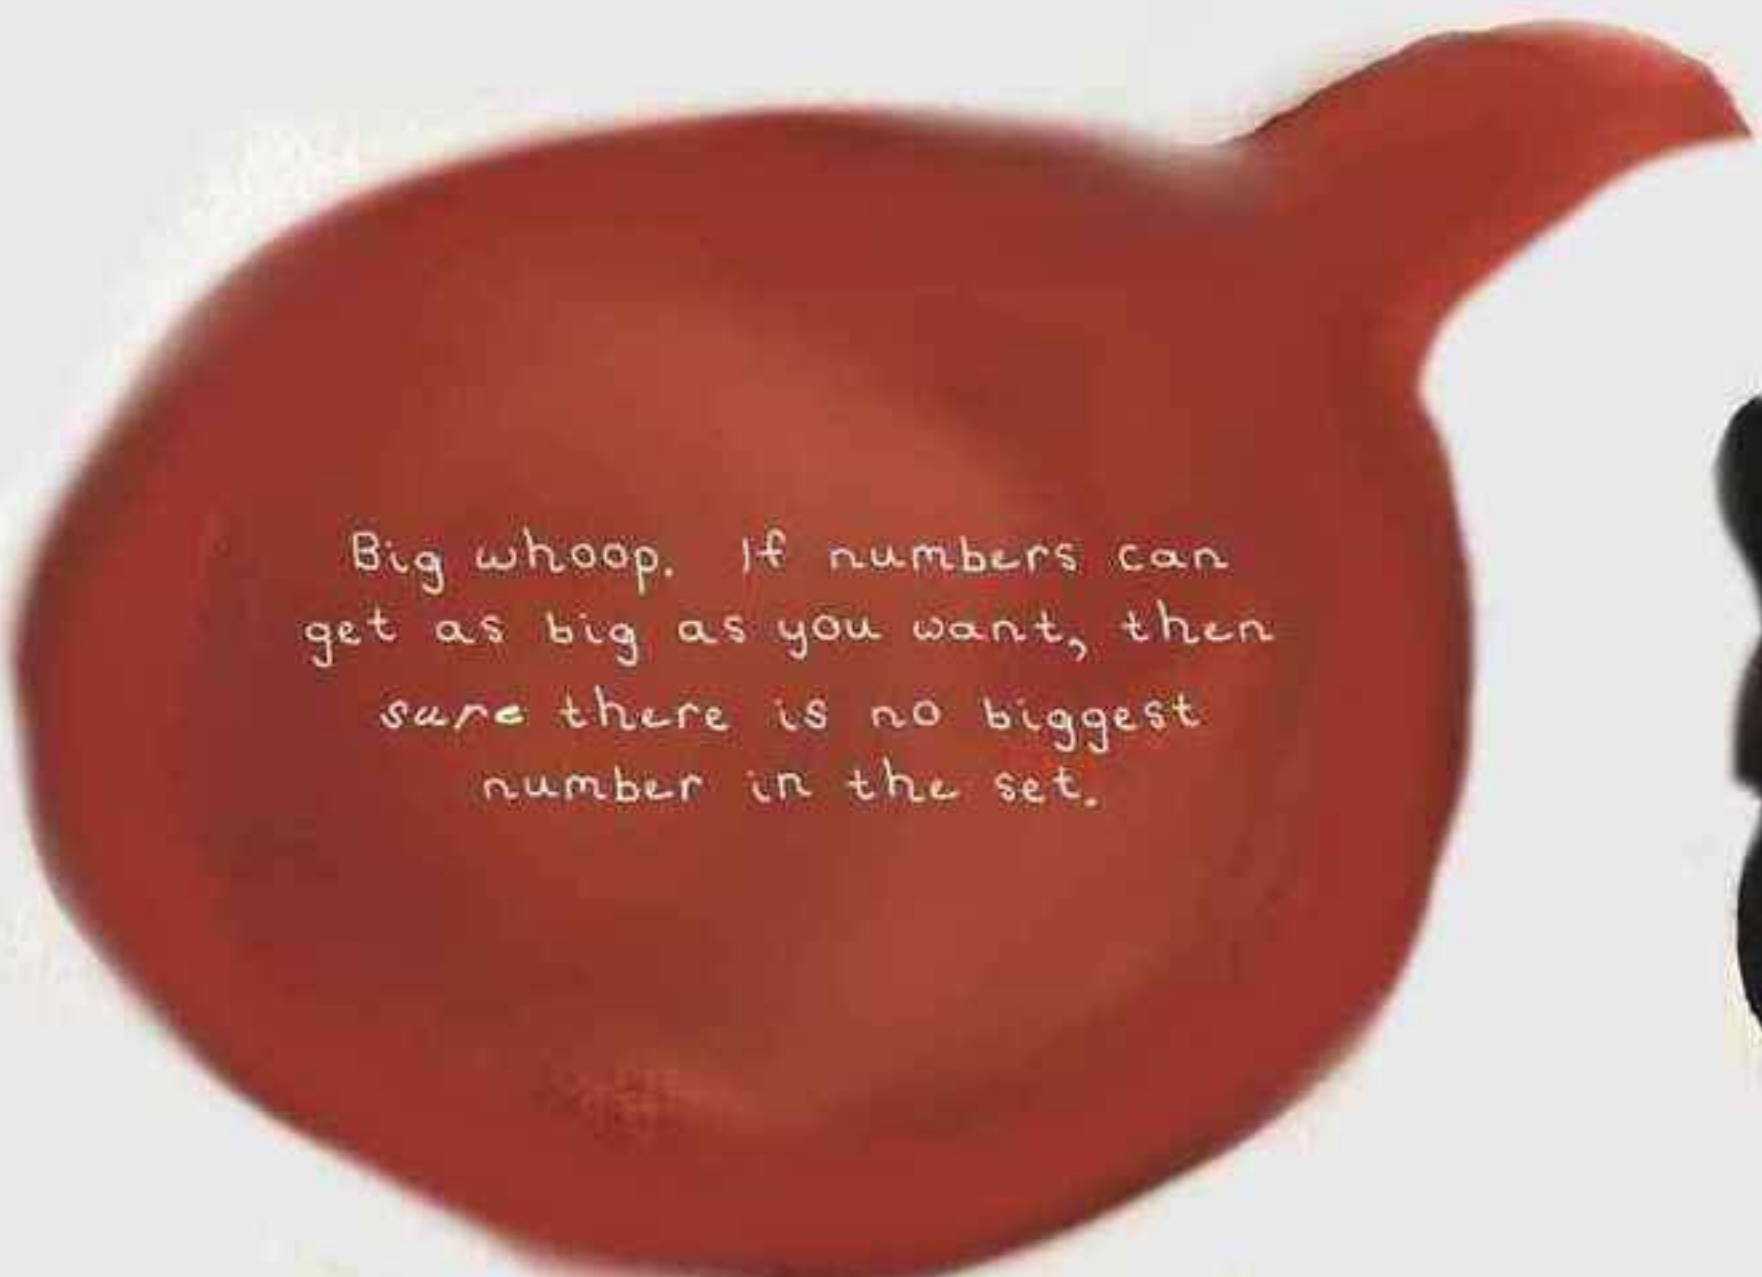
That set (the set of positive integers) has no biggest number.

If I asked you to prove why, what would you say? Well, whenever a mom of coolness x is in the set, so too is the mom with coolness $x+1$, so x cannot be the biggest number in the set. That holds for every single x in the set.






So, some sets don't have
a biggest number.



Big whoop. If numbers can
get as big as you want, then
sure there is no biggest
number in the set.



So, let's make this
more interesting...



But first, a note.

Isn't infinity the biggest element in the set $\{1, 2, 3, \dots\}$?

Why isn't infinity the biggest element in the set of positive integers $\{1, 2, 3, \dots\}$?



Well, infinity isn't an integer. So it can't be in the set of positive integers at all.




Yes, numbers in the set get as high as you'd like them to get. But they don't end - not at infinity or anywhere else.



How do you prove
infinity isn't an integer?

Sure, infinity seems like an
integer at first...

...until we take a closer
look under its disguise.



So let's prove that infinity isn't an integer.



First, let's officially define infinity. One agreeable definition
for infinity is that it's bigger than any other integer.



So let's suppose (for the purposes of contradiction) that infinity is an integer.


Then, what is infinity + 1?



Note that infinity + 1 must be an integer, since an integer plus an integer should equal an integer.


If you say
 $\text{infinity} + 1 = \text{infinity} \dots$

...then, when you subtract infinity from both sides, you get $0 = 1$. That can't be right.



If you say
 $\text{infinity} + 1 > \text{infinity} \dots$

...then that contradicts the fact that infinity was supposed to be larger than every other integer.



If you say
 $\text{infinity} + 1 < \text{infinity} \dots$

...then that's just silly. Adding one to a number should make it bigger, not smaller.



So infinity + 1 can't be equal to infinity, greater than infinity, or less than infinity. So infinity must not be an integer at all.



Does every upper bounded set have a biggest element?

Oops. I got a little sidetracked by these infinity shenanigans. Let's rewind.

I was saying that sometimes sets are unbounded (like $\{1, 2, 3, \dots\}$), and therefore have no biggest element.



Well, fine. But what if the set is bounded?



That is, what if every number in the set is less than or equal to some upper bound we call "alpha"?

Then the set has to have a biggest element, right?



Oddly enough, no. You would think that because the elements of the set are upper bounded, there has to be a biggest element. But consider the set $\{0.9, 0.99, 0.999, \dots\}$.



The elements in the set are all upper-bounded by 1, but there is no biggest element.

And I'll even prove that the set has no biggest element.

Tell me a number in the set that you claim is the biggest...



...and I'll add a 9 to the end, and come up with an even bigger number in the set.



Therefore, any number that you claimed is the biggest cannot be. So we can agree that the set has no biggest number.



What's the most important upper bound of a set?

What's the upper bound of our set $\{0.9, 0.99, 0.999\dots\}$?



I think you'd say the upper bound is 1.



But 1 isn't the only upper bound of that set.

1...



5....



... and 1 billion are all upper bounds of that set.




They are upper bounds because they are all numbers bigger than every number in the set.



But some of those upper bounds seem sort of useless and redundant, like 1 billion.

So let's define the supremum to be the *least* upper bound.



(Which in a sense, is the most important upper bound.)



For example, the supremum of the set $\{0.9, 0.99, 0.999, \dots\}$ is 1. Although we can't know for sure, because we haven't proved it, yet. Just take my word for it, for now.



Does every upper-bounded set contain its supremum?

Does every upper-bounded set have a supremum?

Yes.




Does every upper-bounded set contain its supremum?

No.




For example, the set $\{0.9, 0.99, 0.999\dots\}$ has a supremum of 1, but never actually reaches 1 in the set, and therefore doesn't contain 1.*



*You might say the set does contain 1, since the set goes on forever, and 0.999 repeating forever equals 1. But while the numbers in the set get arbitrarily close to "0.999 repeating forever" (aka 1), they never do reach it.

What's the most important
lower bound of a set?



Does every lower-bounded set contain its infimum?

Ok, so maybe I won't completely spare you.

Does every lower-bounded set have an infimum?

Yes.




Does every lower-bounded set contain its infimum?

No.



For example, the set $\{0.1, 0.01, 0.001, \dots\}$ has an infimum of 0, but doesn't contain 0 in the set.



Alright. That's all the repetitiveness for now, I promise.

What's the point of formal definitions?

We previously said that:

The supremum is the smallest upper bound.

and

The infimum is the biggest lower bound.



But more formally we can say that:

An upper bound ' u ' of set ' S ' is called a supremum of S if for all upper bounds b of S , $u \leq b$.

and

A lower bound ' l ' of set ' S ' is called an infimum of S if for all lower bounds b of S , $l \geq b$.



No kidding? We can define a supremum using a bunch of math words and symbols? Well, isn't that wonderful.

So what's the point of this formal definition?

Sometimes, there is no point (just math people using the weird language they've always used).


But sometimes, formal definitions come in handy during proofs.

For example, say I wanted to convince someone that the infimum of $\{.1, .01, .001, \dots\}$ is 0.

What if that someone insisted the infimum is 0.0036?



I'd never convince them if my only reason for why 0 is the greatest lower bound is: "Uhh...it's definitely a lower bound...and it's definitely the biggest one..believe me."



In order to *actually* convince someone, I'd have to go through a proof by contradiction, and I'd need to use the formal definition. Let's do that next.




How can we prove something is an infimum?

Let's prove 0 is the greatest lower bound of the set $\{0.1, 0.01, 0.001, \dots\}$.

First, let's prove that 0 is a lower bound.

Well, yes, 0 is smaller than any element of the set.




Finally, let's prove that 0 is the greatest lower bound.

Let's do so by contradiction.

Let's assume that there was a number strictly bigger than 0 (call it alpha) that was the actual greatest lower bound.



If we replace the two most significant digits of alpha with "01" and truncate everything after that... then we get an element in the set that is definitely smaller than alpha.



So alpha can't be a lower bound of the set (because we found something in the set smaller than it)..

Thus, 0 must be the greatest lower bound.



What's the point of formal proofs when we have intuition?

Our intuition would tell us that the greatest lower bound of $\{0.1, 0.01, 0.001, \dots\}$ is 0.




And it would be right.




But if our gut intuition can do the work of proving things, then what's the point of formal definitions and proofs?



Well, it might be useful if mathematicians we worked with came up with the same results we did.



Despite different intuitions, we want consistent results.



Is $\inf(A) = -\sup(-A)$?

Now, for an even more challenging proof.

Say set A is any set of real numbers with a lower bound.
Say set $-A$ is the set of the negatives of those numbers.
Let's prove $\inf(A) = -\sup(-A)$.



First, let's consider an example.

Say $A = \{1, 2, 3, \dots\}$.
Then $\inf(A) = 1$.



Now let's flip it.
Then $-A = \{-1, -2, -3, \dots\}$.
And so $-\sup(-A) = -(-1) = 1$.



So the equation checks out, at least for this particular set A .


Now, let's prove the equation holds for every set ever.
We'll do so in two steps...



In this step, we will prove $\inf(A) \leq -\sup(-A)$.
In the next step, we will prove $\inf(A) \geq -\sup(-A)$.
If both are true, that must mean $\inf(A) = -\sup(-A)$.

Step 1: Prove that $\inf(A) \leq -\sup(-A)$.

We know A has an **infimum**,
since it's bounded below.



And by the definition of
infimums:

$\inf(A) \leq a$ for any a in A .



Let's flip the inequality,
multiplying both sides by -1 :

So **$-\inf(A) \geq -a$** .



But guess what? The element
" $-a$ " could be any element in
 $-A$.

So **$-\inf(A)$** upper bounds $-A$.



So since $-\inf(A)$ is an upper bound of $-A$...
it has to be bigger than the least upper bound of $-A$:

$-\inf(A) \geq \sup(-A)$



So **$\inf(A) \leq -\sup(-A)$** .

Step 2: Prove that $\inf(A) \geq -\sup(-A)$.

We know $-A$ has a **supremum**, since it's bounded above.




And by the definition of supremums:

$\sup(-A) \geq -a$ for any $-a$ in $-A$.



Let's flip the inequality, multiplying both sides by -1 :

$-\sup(-A) \leq a$.



But guess what? The element "a" could be any element in A .

So **$-\sup(-A)$** lower bounds A .



So since $-\sup(-A)$ is a lower bound of A ...
it has to be smaller than the greatest lower bound of A :

$-\sup(-A) \leq \inf(A)$.



So $\inf(A) \geq -\sup(-A)$.

In step 1, we proved $\inf(A) \leq -\sup(-A)$.
In step 2, we proved $\inf(A) \geq -\sup(-A)$.
Since both are true, that must mean **$\inf(A) = -\sup(-A)$** .

What is a number?

If someone math-splained what a number was to me, I'd probably want to slap them in the face. But as it turns out, trying to define a number is more difficult than I thought it would be.


Is it something that counts items?



Then what of π ?



Is it something that measures lengths?



Then what of $\sqrt{-1}$?



Let's take a step back. What do we do to numbers?




We add and multiply them.



So let's say numbers are the things we can add and multiply...

...and certain properties of addition and multiplication hold.

Which properties should hold? Well, it would mean trouble if $1+2$ didn't equal $2+1$. There are a bunch more troublesome scenarios, too.



So, in order to avoid all trouble, let's establish some rules. Let's say...

Addition should be commutative.
 $a+b = b+a$

Addition should be associative.
 $(a+b)+c = a+(b+c)$

Multiplication should be commutative.
 $ab = ba$

Multiplication should be associative.
 $(ab)c = a(bc)$

Addition and multiplication should work together via the distributive property.

$$a(b+c) = ab+ac$$

So let's say a number is anything that behaves according to those rules.



What is a complex number?

Yes, this is *real* analysis, so we won't usually deal with *complex* numbers, but let's go ahead and define one.

The vanilla way to define a complex number is to say it's the sum of a real and imaginary number, like $a+bi$.



But let's do it a more whimsical way. Let's say each complex number is an ordered pair, like (a,b) .



Let's add them like:
 $(a,b) + (c,d) = (a+b, c+d)$




Just for kicks, let's multiply them like:
 $(a,b) \times (c,d) = (ac-bd, ad+bc)$



And instead of the usual $i = \sqrt{-1}$, why don't we say:
 $i = (0,1)$



Bam, I've defined the complex numbers.



I seem to have just told you two different definitions of complex numbers, which seem to have nothing to do with each other. We can write them as either of the following:

$$a + bi, \text{ where } i = \sqrt{-1}$$
$$(a, b), \text{ where } i = (0, 1)$$



But, it happens that the two definitions are equivalent.

For example, if we calculate $i \times i$ the vanilla way, we get $a = -1$ and $b = 0$.

$$\begin{aligned} i \times i &= \sqrt{-1} \times \sqrt{-1} \\ &= -1 \\ &= -1 + 0i \end{aligned}$$



And if we calculate $i \times i$ the whimsical way, we still get $a = -1$ and $b = 0$.

$$\begin{aligned} i \times i &= (0, 1) \times (0, 1) \\ &= (0 \times 0 - 1 \times 1, 0 \times 1 + 1 \times 0) \\ &= (-1, 0) \end{aligned}$$



What should a universe of numbers look like?

Some sets of numbers seem to act more like a self-sufficient "universe" than other sets. For example...

The set of real numbers seems to act like its own universe of numbers. Because no matter how you add, subtract, multiply, or divide two real numbers, you always end up with a real number.



But that's not true for the set of integers. For example, if you divide the integer 1 by the integer 2, you end up at $1/2$ which is outside the universe of integers.



Let's call a universe of numbers a *field*, and let's define it to be any set of numbers that adheres to the following three properties:

When x and y are in the field, then the field must also contain...

... $x+y$ and $x*y$.

(Closure)



*So the set $\{1,2\}$ is not a field, because $1+2$ is not in the set.

...the additive identity (0) and multiplicative identity (1).

(Identity)



*So the set of positive integers $\{1,2,3,\dots\}$ is not a field, because 0 is not in the set.

...the additive inverse ($-x$) and the multiplicative inverse ($1/x$) when x is nonzero.

(Inverses)




*So the set of integers is not a field, because 2 is in the set, but its multiplicative inverse $1/2$ is not.

And so, no matter how you add, subtract, divide, or multiply numbers in a field, you stay in the field - its own inescapable universe.

It's useful to work with fields, because a bunch of theorems only hold within fields.

That is, in fields, numbers work in a way we expect.



But when we're not dealing with a field, a bunch of theorems don't hold, and all sorts of crazy things can happen.*



*To be honest, it's not like everything is always chaos outside of a field. There exist quite useful and orderly structures that are not fields, like rings and groups (which come up a bunch in abstract algebra, though not so much in real analysis).

How should we order fields?

Sure we could order *fields* like we order sets. That is, we could say that a field is any set of numbers with an ordering imposed on it (i.e., a way to determine which numbers are bigger and smaller).



But there's a problem. Let's try to order the complex field like that. Our order can be that $(a,b) < (c,d)$ when $a < c$, or when $a=c$ and $b < d$. For example, $(1,999) < (999,1)$ because $a < c$. And $(2,3) < (2,4)$ because $a=c$ and $b < d$.



We end up able to prove contradictory results (like that 1 is both less than and greater than -1).

Proof:

We know $-1 < 1$, since by our definition, $(-1,0) < (1,0)$.

We know i is positive, since by our definition, $(0,0) < (0,1)$.

So $-i < i$, since we can multiply by a positive without flipping the inequality.

So $-i < i$, by multiplying both sides by i again.

So $1 < -1$.

Contradiction!




The flaw in the proof is that our definition of an ordered field doesn't ensure that if $a < c$, then $ba < bc$ (when b is positive).

We probably want an ordered field to ensure that. Because if a field is ordered, we should be able to compare numbers with inequalities to determine which order they go in. But with this definition, we can't even order 1 and -1.

So let's create a stricter definition of 'ordered field', one that will let us work with inequalities in a more sensible manner.

Let's define an ordered field to be a field in which:

- 1) Each number in the field is either positive, negative, or 0.
- 2) If a and b are positive, then $a + b$ is positive
- 3) If a and b are positive, then $a * b$ is positive.



Under this definition, complex numbers are not an ordered field (thank goodness, because intuitively they shouldn't be, since we couldn't find out how to order i and $-i$). Let's prove it using the definition.

Is the complex field an ordered field?

According to our definition of 'ordered field,' each number in an ordered field must be either positive, negative or zero.

So if the complex field is an ordered field, then i (aka $\sqrt{-1}$) must be either positive, negative, or zero.

We know i can't
be zero...

Because we already
have a number that's
zero - it's called
zero.

And i can't be
positive...

Because since
 $\text{positive} * \text{positive} = \text{positive}$
by our definition of an
ordered field, i^2 should
be positive.

But $i^2 = -1$, which is
negative.

And i can't be
negative...


Because since
 $\text{positive} * \text{positive} = \text{positive}$
by our definition of an
ordered field, $(-i)^2$
should be positive.

But $(-i)^2 = -1$, which is
negative.

So the complex numbers can't make up an ordered field,
because they violate the first axiom of ordered fields - that
each number in the field is either positive, negative, or zero.

How do we construct
the real numbers?

Why is "constructing" the real numbers (that is,
comprehensively defining and building up the real numbers)
such a big deal in analysis?



Why build them up when we already intuitively know what they
are? I've never received a satisfying answer for this.



The best answer I can come up with is "because we can."

So how to do it? How to define the real numbers?

Apparently we can't do it by drawing a real number line and saying anything on it is a real number...

...because that process is not well defined. "Real number line" means nothing without defining "real number."



And apparently we can't do it by saying anything that can be written as a decimal is a real number...

...because we'd have to prove that we can add and multiply real numbers...


..but how can we add numbers with an infinite number of decimals? There is no rightmost digit to start from.



$$\begin{array}{r}
 3.14158392121 \\
 + 2.71827005712 \\
 \hline
 ???
 \end{array}$$


So let's use a technique called "Dedekind cuts."

It's a method of constructing the real numbers from sets of rational numbers.





Isn't that neat? We can take something smaller (the rationals) and use it to create something bigger (the reals).

First, to represent each real number α , we create one subset (or "cut") of rational numbers.



Specifically, cut α must be nonempty, contain all rationals less than α , and have no greatest element.

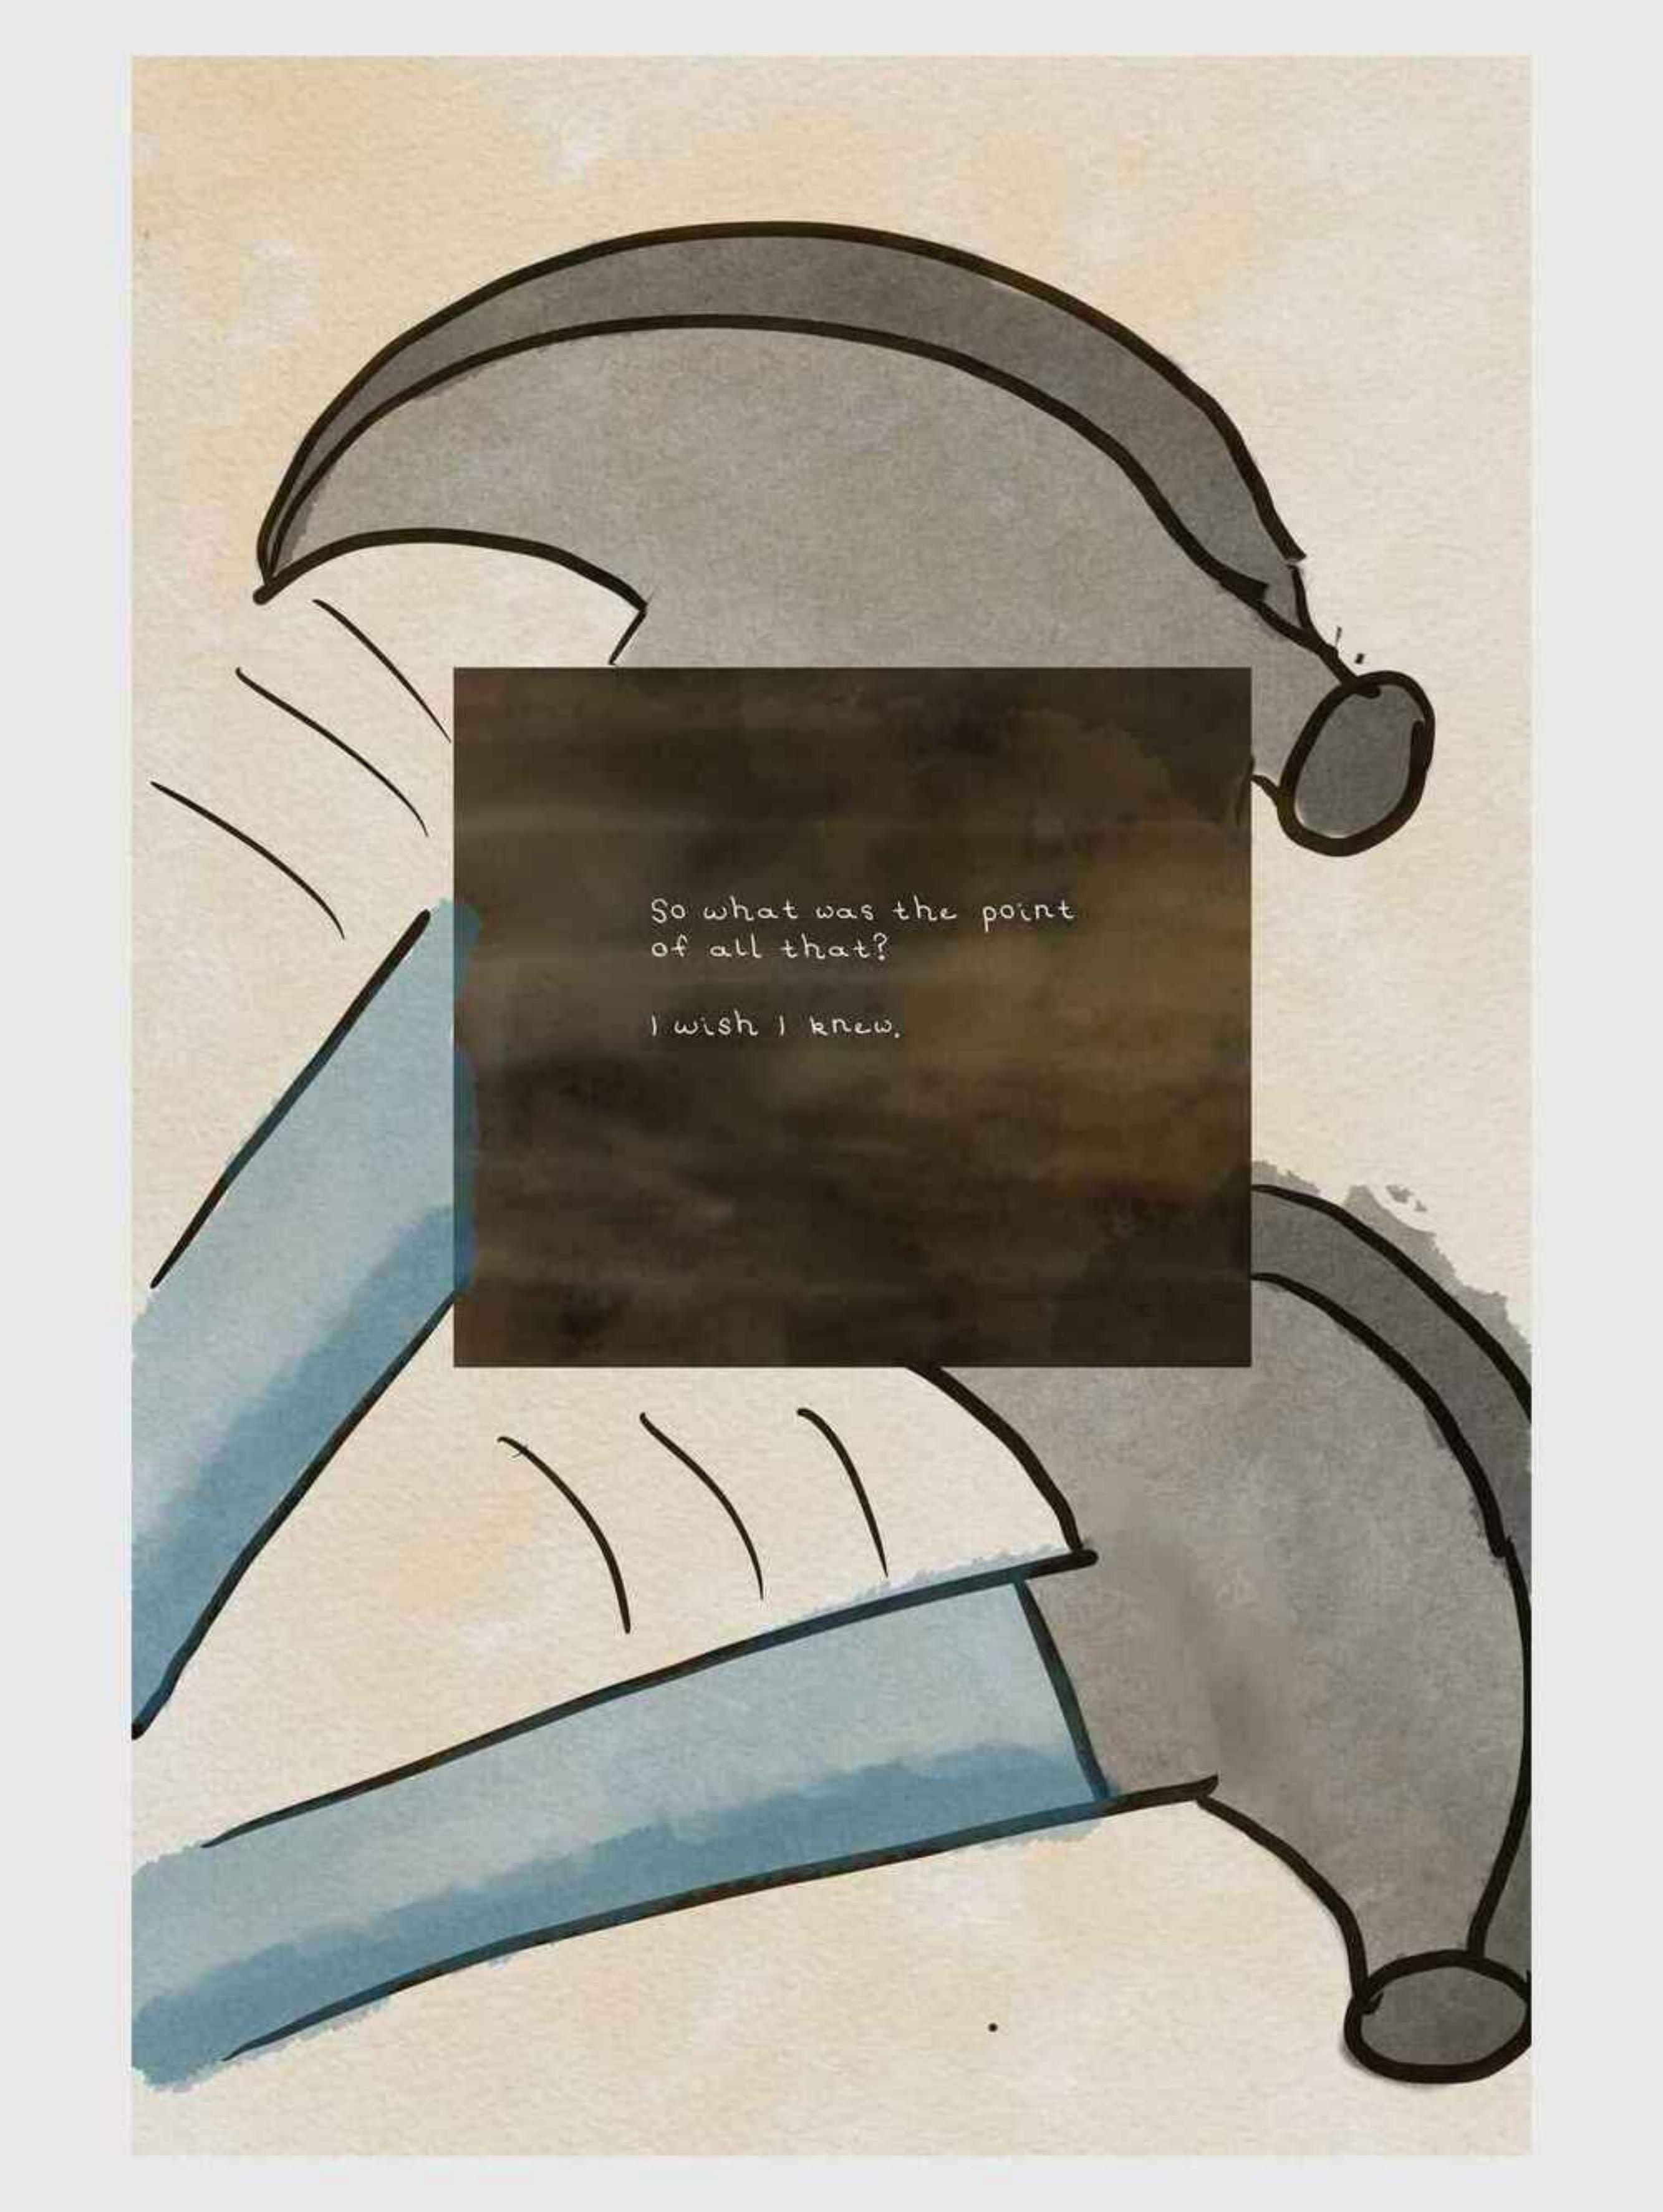
For example, the cut "pi" is the set that includes 2.5, 3, 3.141, and every other rational number less than pi.



The final step is to prove every real number can be represented by one and only one of these "cuts." That part involves showing we can order, add, and multiply these cuts. But it gets a bit tedious, so I won't do it here.



But once we've done that, we've constructed the reals from the rationals.



So what was the point
of all that?

I wish I knew.

(Brainteaser) Are irrationals
a subset of the reals?

I had thought so...

There is $\sqrt{2}$... $\sqrt{3}$... π ...



But actually, no!

The imaginary number $i = \sqrt{-1}$
is also irrational.



Are there infinitesimally small elements in the real line?

No. Real numbers cannot be infinitesimally small, according to the Archimedean property of the reals. It says this.

Find any real number. Call it ϵ (epsilon). It can be really small.



Then, there exists a natural number n such that $1/n$ is even smaller than that ϵ .



*There's always someone smaller than you
Yeah you're the smallest right now
But you won't be it for long
There's always someone smaller than you...*

I used to think the Archimedean property of the reals was obvious - of course you can always find a number smaller than another!

Then I realized that dx and dy (infinitesimals) do not satisfy the Archimedean property.



So the Archimedean property really is a property somewhat unique to the reals.



So, there are no infinitesimally small elements in the real line. You can always find a real smaller than another real.

How can we raise a number to a negative power?

Do real numbers make sense in exponents? For example, what does a negative exponent mean?



Sure, you might have memorized that a^{-c} is the same as $\frac{1}{a^c}$.



But does that even make sense?



Exponents are supposed to mean "multiply the thing on the bottom by itself as many times as this exponent tells you to."

$$2^3 = 2 \times 2 \times 2$$

So then how do we multiply something by itself negative times?

$$2^{-3} =$$



Well we can prove* that for positive exponents, the following two properties hold:

$$a^b \times a^c = a^{b+c}$$

$$\frac{a^b}{a^c} = a^{b-c}$$

So, if we want that summing rule to still work with negatives exponents, then a^{-c} must equal $\frac{1}{a^c}$.

The diagram shows a derivation of the property for negative exponents. It starts with the summing rule $a^b \times a^c = a^{b+c}$ from the previous panel. An arrow points to the equation $\frac{a^b}{a^c} = a^{b-c}$. A note "rule 1.10" points to the exponent $b-c$. Then, an arrow points to the equation $\frac{1}{a^c} = a^{-c}$, with a note "simplify" pointing to the exponent $-c$.

*Although I won't prove it, out of laziness, mind you.

Similarly, what does it even mean to raise something to an irrational power?



Maybe it's obvious to you, but it shouldn't be. It's not something that's inherently true.



People defined it that way, and then taught it to us that way, and now we just take it as truth.



To simplify all sorts of future conundrums, let's just say the following property always holds:

$$a^b \times a^c = a^{b+c}$$

$$\frac{a^b}{a^c} = a^{b-c}$$



...no matter how much nonsense it is to raise numbers to negative or irrational powers.




What happens when we put a bunch of real lines together?


We can mash together real lines to get higher dimensional spaces.




If we have just one real line, we'll call that space \mathbb{R}^1 .



If we have two real lines, we have a real plane: \mathbb{R}^2 .



If we have three, we get the space called \mathbb{R}^3 . And so on.



When we put any n such real lines together, we get euclidean n -space (\mathbb{R}^n). To be more formal, euclidean n -space is a vector space where the points (x_1, \dots, x_n) contain all real numbers.



What happens to the reals
when we allow
dividing by zero?

Yes, you're not "supposed" to divide by zero. But what happens when you allow dividing by zero in the reals?



Well, you get $-\infty$
(when you divide a
negative by 0) and
 $+\infty$ (when you divide
a positive by 0)...



...which gives us the extended reals!
The extended reals are just the reals,
plus the infinity symbols $-\infty$ and $+\infty$.



What's so special about
the reals?

Just like how certain theorems only hold for *fields*, certain theorems only hold for the real number field.

In fact, so many theorems hold for *just* the reals that if you're enough of a chump, you could write a whole book about them.



Conclusion

So that's it for this first chapter: an introduction to the real and complex number systems.

A fun problem you could try now that you know this stuff:

- Show that if the additive identity has a multiplicative inverse, then the field must have only one element.

If you want to solve some related open problems that no mathematician has ever solved before:

- Solve the Integer Chebyshev problem (minimize the supremum of the Chebyshev polynomial).
- Prove the Euler-Mascheroni constant is irrational.

Enjoy!



Real Analysis

Chapter 2
Infinities

Figuring out how to deal with infinities is especially important for proving calculus works. After all, I'm not at all convinced (and you shouldn't be either) that you can just sum something infinitely, and end up with a finite number, like we often do when we integrate.

And part of figuring out how infinities work, is, of course, comparing them.

How do you compare
infinities?

How can we even begin to determine...

...whether some infinities...




...are bigger than others?



I guess we could just compare infinities using the same technique we use to compare everything else...that is, by counting up each and seeing which is bigger.


Even though we know we compare finite quantities by "counting" them up and seeing which set is bigger, the tricky part is coming up with a formal definition of "counting."

Is this counting?




No, because I didn't count everything.

Fixed! Now, is this "counting"?




No, because I double-counted.

So, real "counting" must be pairing each object to exactly one number - no more, and no less.




A less wordy way to say "pairing each object in one set to exactly one object in the other" is to say "bijection."

Formally, a bijection is a function that maps every element of a set X to a unique element in set Y , and maps every element in set Y to a unique element in set X .



When we "count," what we're actually doing is creating a bijection between the set of numbers we are counting with...



...and the set of things we are counting.

In general, whenever we create this bijection, or one-to-one pairing, between sets, we know they are the same size.



So this is how we count finite quantities - by constructing bijections. We just never really use that terminology.

And, as it turns out, we can use these same bijections to compare infinities as well.

Which set is bigger:
the integers or even integers?

It seems like there should be more integers than even integers (after all, the set of even integers are contained within the set of integers). So let's just verify. According to our definition...



...if there's a way to pair two sets so each item in one set is paired to exactly one item in the other set...



...then the sets are exactly the same size.



And so, because we can biject the integers to the evens using the pairing rule " n maps to $2n$ " (for example, 1 maps to 2, 2 maps to 4, and so on...) the sets are the same size.

Of course, this seems completely ridiculous.

The even integers are a subset of the integers, and yet we get that they are the same size.

These are the sort of absurdities that arise when we try to do something as absurd as counting infinities, I guess.

Galileo actually presented this weird result as a paradox to say we shouldn't be comparing sizes of infinity at all.

How do we resolve that paradox? Do we just agree with Galileo and stop counting infinities?

Well maybe it's not the end of the world if some of our intuitions about infinite sets (e.g. that subsets should be smaller than supersets) don't extend to infinite sets (where apparently subsets can be the same size as supersets).

Of course, feel free to come up with your own ways of measuring infinity. We all know, "that's just the way it's done" doesn't really mean it can't change eventually.

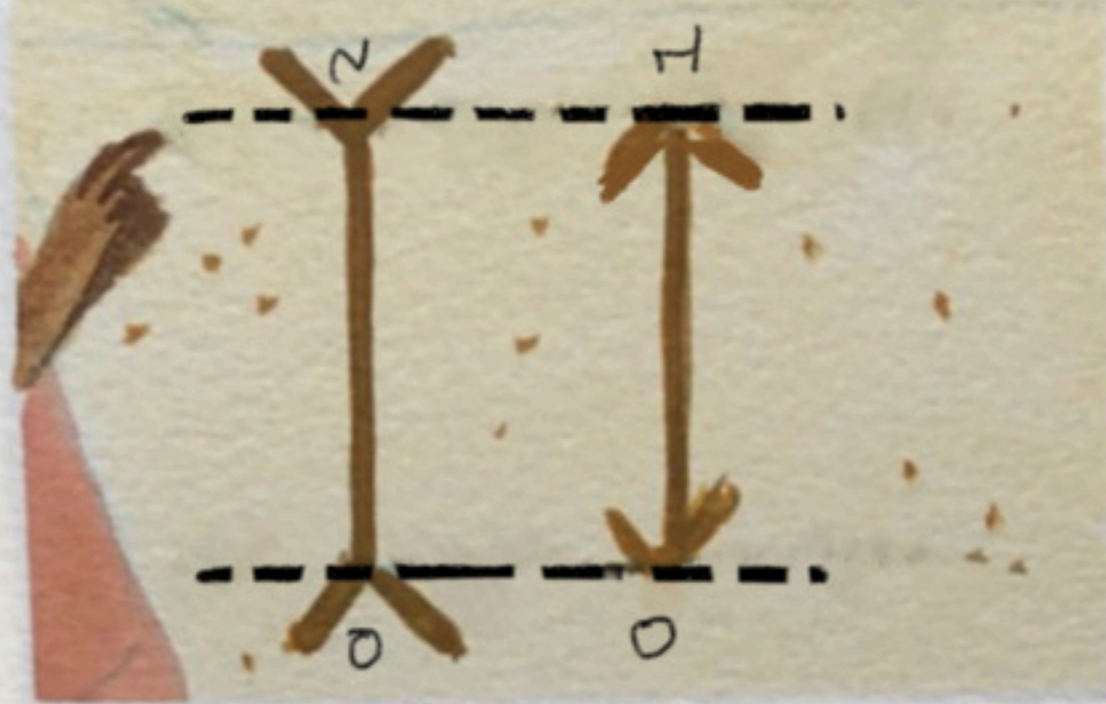

Which set is bigger:
 $(0,1)$ or $(0,2)$?

Hazel, from John Green's *Fault in Our Stars* says:

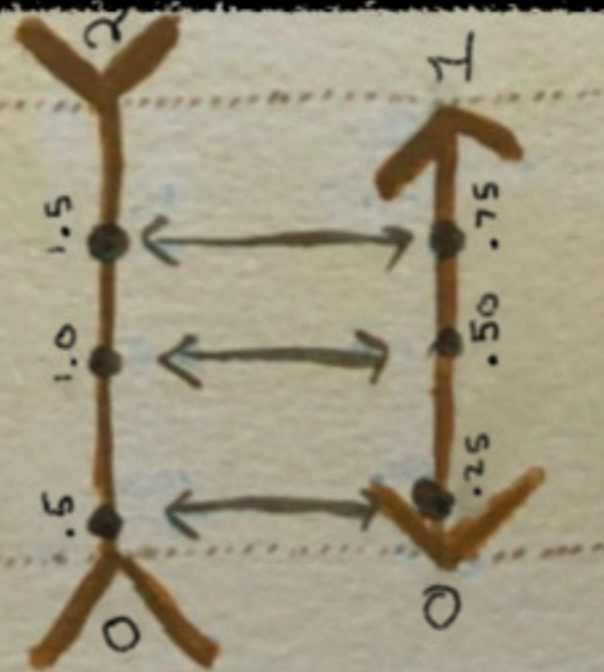
"There are infinite numbers between 0 and 1...Of course, there is a bigger infinite set of numbers between 0 and 2."

In a few pages, we will prove part of what Hazel said right — that some infinities are bigger than others. But just as a fun fact...

...according to our definition, the set of numbers between 0 and 2 is the *exact same size* as the set of numbers between 0 and 1.



We see this by setting up the same bijection as we did before: we pair every number " n " in $(0,1)$ to the number " $2n$ " in $(0,2)$. Because each number in each set is paired to exactly one number in the other set, this is a bijection.



Note: John Green knew Hazel was "incorrect." Here is what he says: "The idea there was that I liked that 16-year-olds could make — as they do — incorrect abstract conclusions about complex mathematics. But even if these conclusions are incorrect, they can provide real and lasting consolation. I felt like it would be too neat/tidy to have everything be correct; I wanted her to make incorrect inferences...that still guide her thinking in a correct/helpful direction."

I think at this point, to be fair to Hazel, I should mention that using bijections is just one way mathematicians measure infinities.

Another method — called the “inclusion order” — implies subsets are smaller than supersets, and would make Hazel correct.


However, it does turn out that using the “bijection” method of counting infinities is what’s most relevant to us in proving calculus works.

Which set is bigger: the rationals or the irrationals?

Are there more rational numbers in
the universe...



...or more irrationals?



I had thought for sure there were more rationals, just because I could list them off more easily, and the only irrationals I knew were pi and e. But as it turns out, there are actually way more irrationals.

We know this because while we can prove that the rationals are the same size as the counting numbers...

(and so we call the rationals a "countable infinity"...))

...the set of irrational numbers is so big that when we can pair every counting number with an irrational, and still have irrationals left over.

(So we call the irrationals an "uncountable infinity."))



But now let's actually prove this is true.

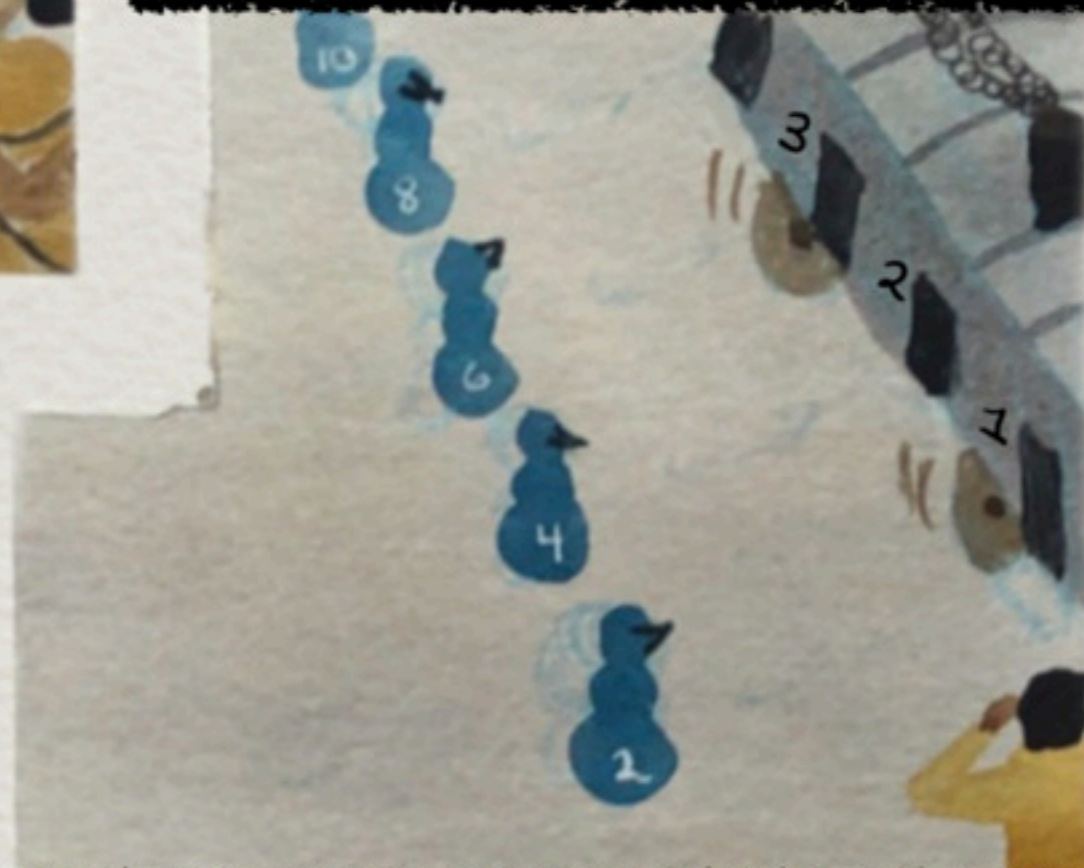
Are the rationals countable?

To see if a set is a countable infinity...

...just pair each element to a counting number, creating a bijection.



For example, we can biject the set of even numbers $\{2, 4, 6, 8, 10, \dots\}$...

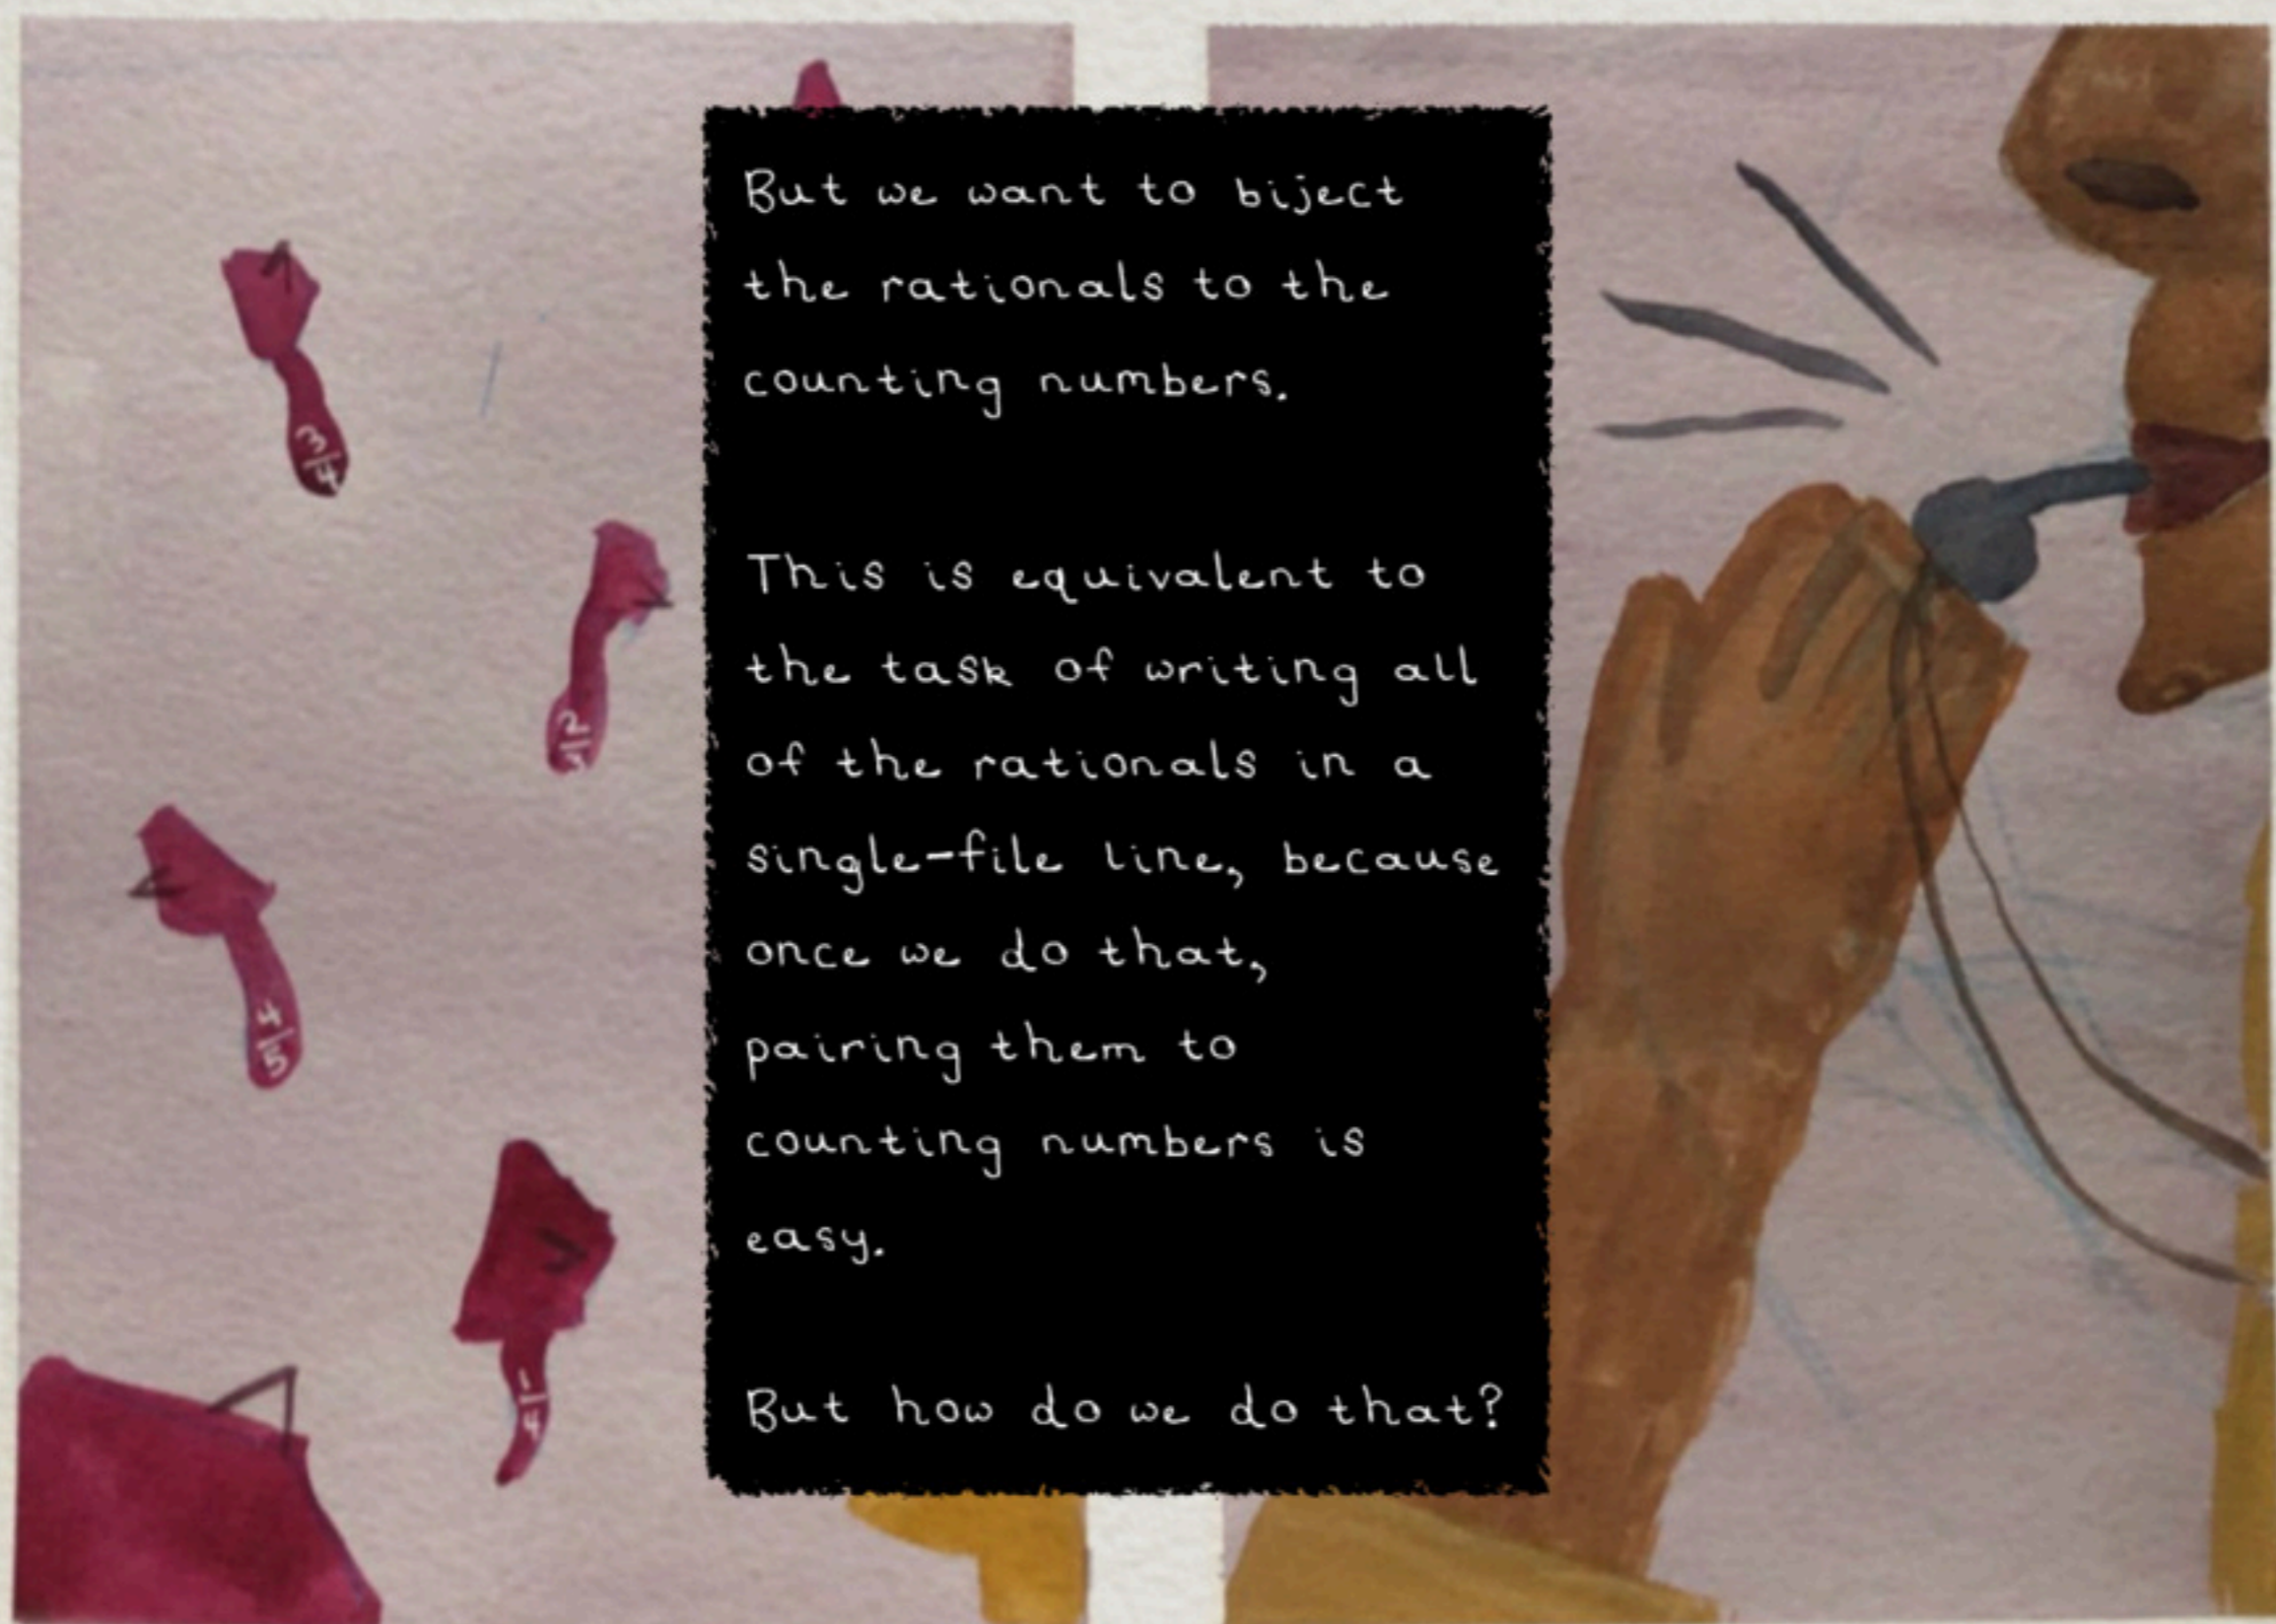


...to the set of counting numbers $\{1, 2, 3, 4, 5, \dots\}$. So, we see the set of even numbers is countable.



But we want to biject the rationals to the counting numbers.

This is equivalent to the task of writing all of the rationals in a single-file line, because once we do that, pairing them to counting numbers is easy.

But how do we do that?



Well, you can't stop me from writing them out in a neverending box like this...



And then drawing a continuous "line" that snakes around them like this.



And so, the rationals can be bijected with the counting numbers, and is therefore a countable infinity.

Are the irrationals uncountable?

Let's prove the irrationals are uncountable (that is, they are so big that they can't be bijected to the counting numbers).

Proof by contradiction. Suppose the irrationals were countable. Then, somebody would be able to give you a list of them, with each one paired with a counting number.



Irrational #1: 3.196254...

Irrational #2: 0.000313...

Irrational #3: 6.913934...

Irrational #4: 8.9102395...

Irrational #5: 0.01016902...

Irrational #6: 2.111357734...

Irrational #7: 6.781122611...

But, no matter what list you give me, I can find an irrational number that is nowhere on that list (and when I do, we have a contradiction.)

No matter what the list you gave me said...

Irrational #1: 3 1 9 6 2 5 4 ...

Irrational #2: 0.0 0 0 3 1 3 ...

Irrational #3: 6.9 8 4 3 4 ...

Irrational #4: 8.7 1 0 2 3 4 5 ...

Irrational #5: 0.0 1 0 1 6 9 0 2 ...

I would take the i^{th} number of the i^{th} digit.

If the number was a 1, I would put a 0 in the i^{th} place of my new number.

If the number was a 0, I would put a 1 in that place.

My number: 1 . 1 0 1 ...

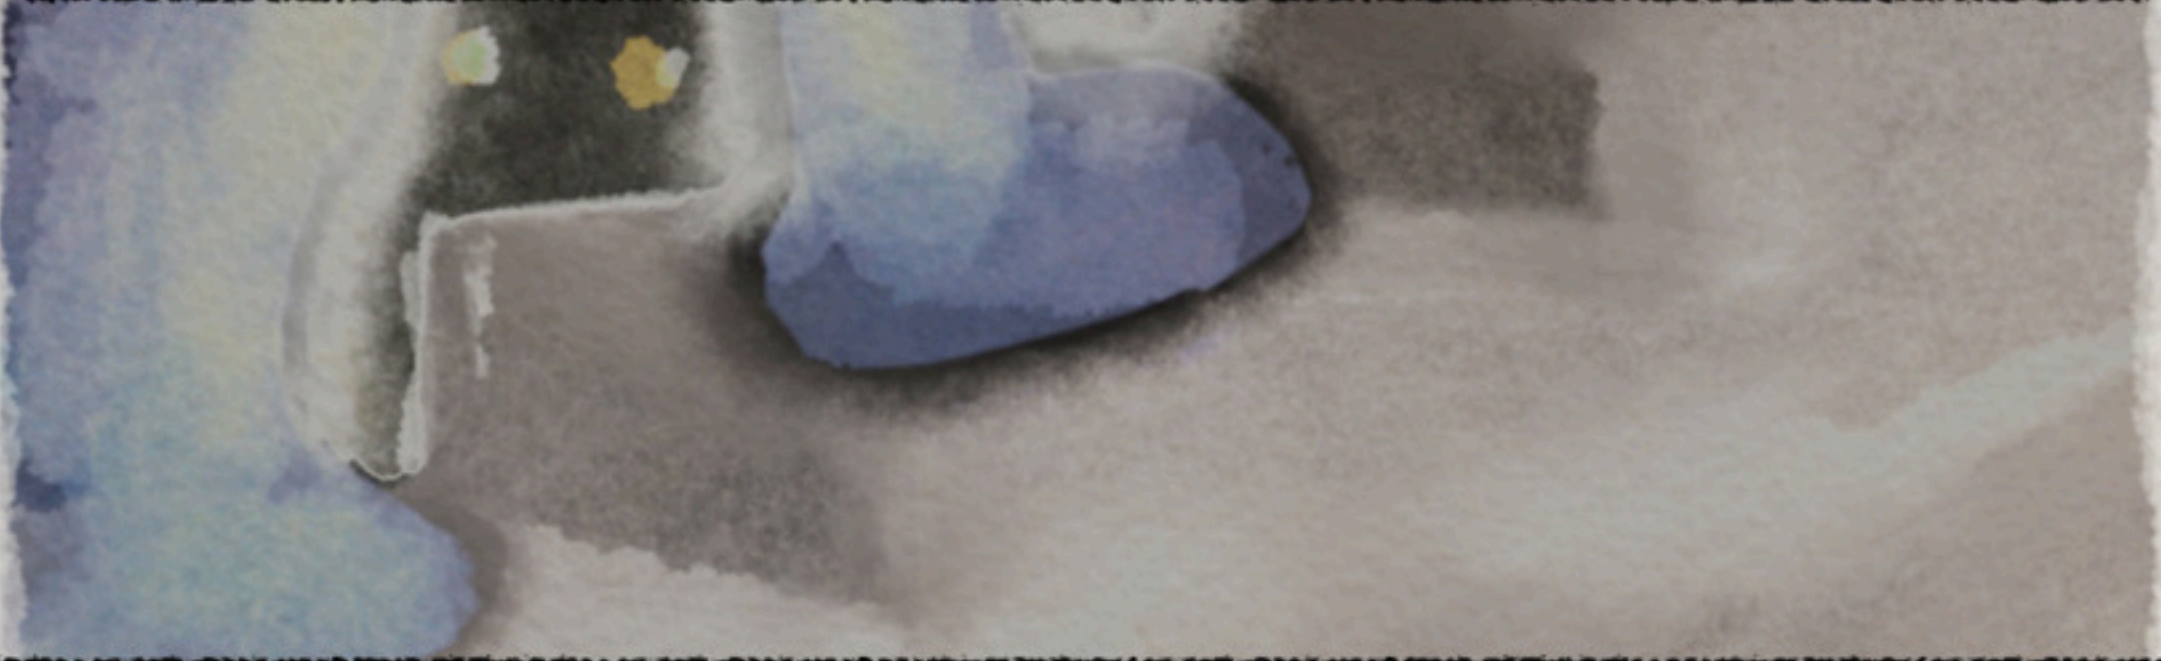
...and in that way, I would end up with an irrational number nowhere on that original list.

Since I could do this no matter what list of irrationals you throw at me, you could never have a complete list pairing each counting number to an irrational number. So, the irrationals aren't countable.

Are there more types of infinity?

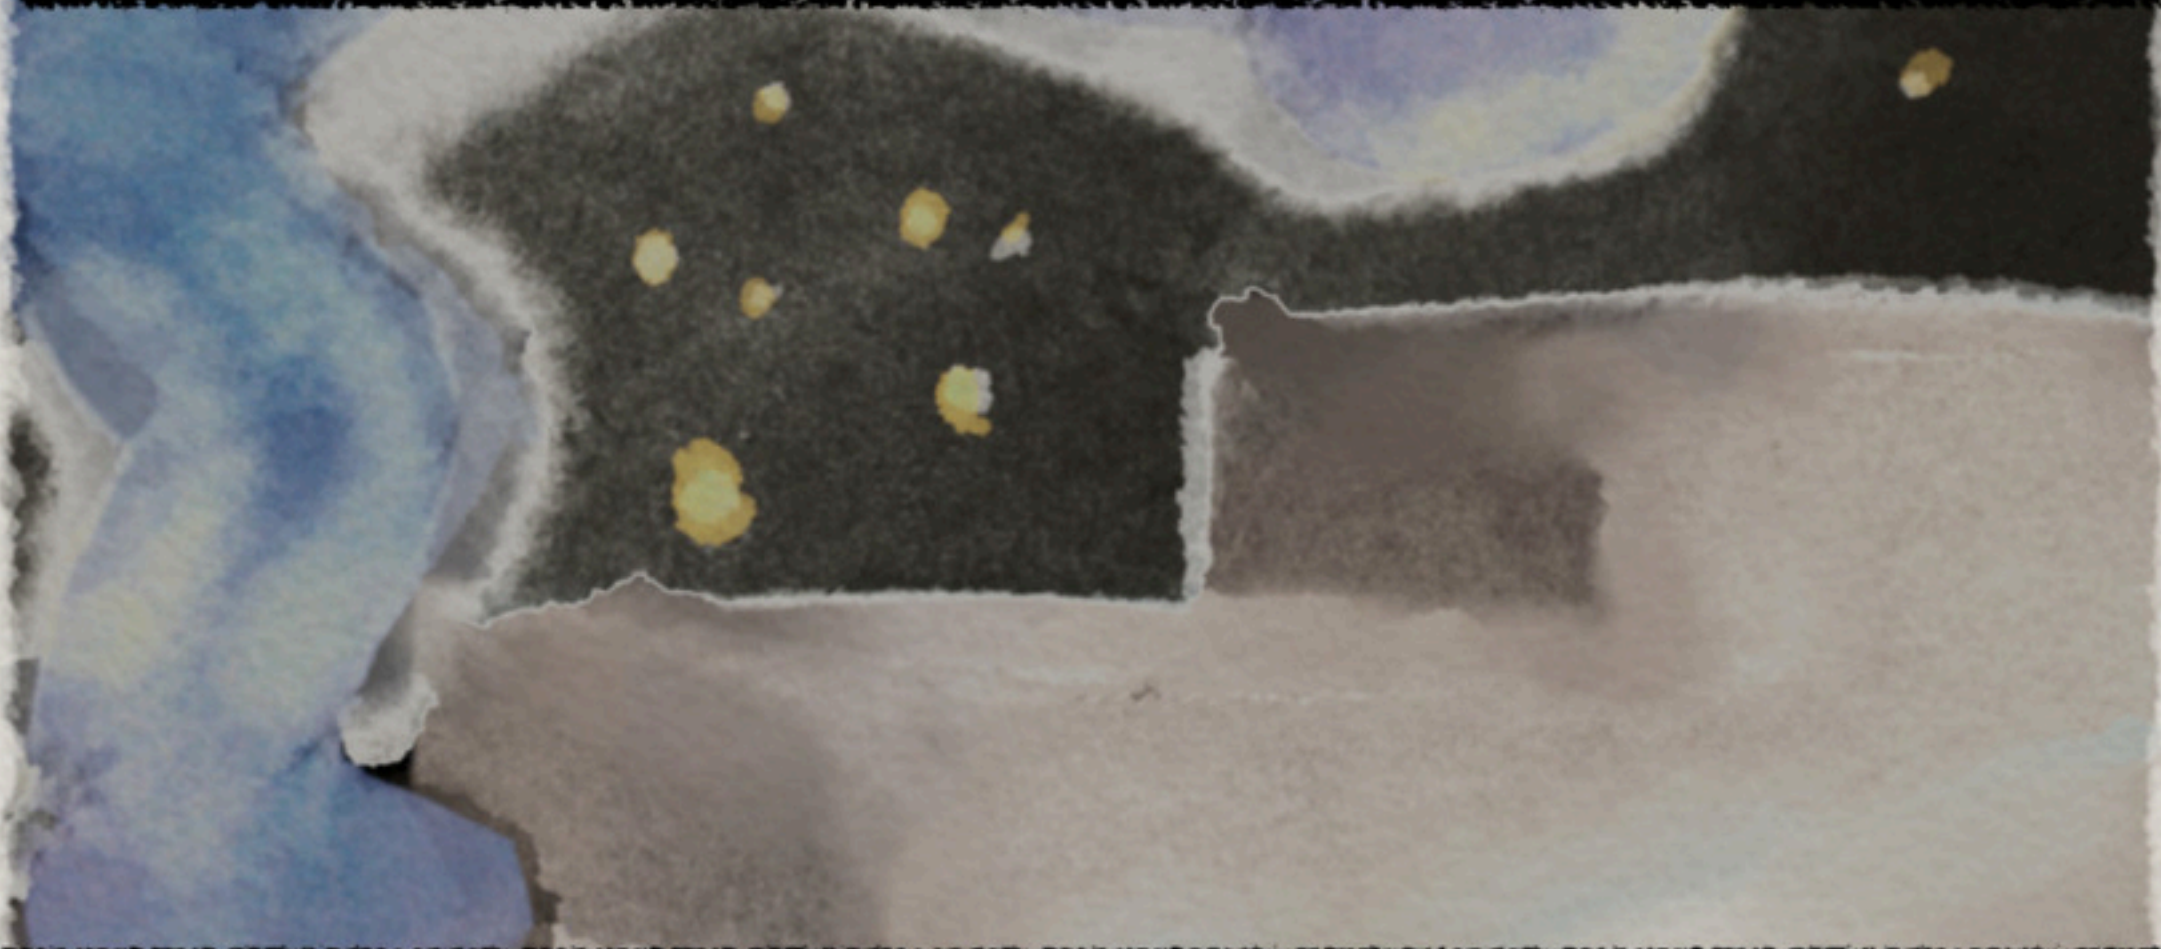
So, we have that the rationals are countable. And the irrationals are an even bigger infinity: uncountable.

But is there a level of infinity in between countable and uncountable?



Well, that question was proved to be "undecidable." What does that mean? It can be taken as true or false without violating any other math (specifically, the widely-used Zermelo-Fraenkel set theory axioms).

But apparently, it is very much decided that you can have an infinity greater than both countable and uncountable.



It's the size of the "power set" of any uncountable set. A power set is the set of all subsets of the set (including the empty set, and the set itself). And we can prove you can never create a 1-1 map from a set to its power set. So that means, if we create a power set of any uncountable set, its size is greater than uncountably infinite.

Brain teaser

Why does the same proof technique we used to show the irrationals are uncountable not work with the rational numbers?

That is, for any list of rationals you give me, why can't I use the same technique to go down the list and find a number nowhere on that list?



Hint:

What properties hold in that new number? What groups can we say that new number belongs to?

Answer:

The same proof doesn't hold, because there's no guarantee that new number you'd find would be rational.

Super-hard brainteaser

For an extra-hard challenge, here's a very hard problem that one of the most clever math majors I knew at MIT couldn't figure out.

For each real number a in the interval $[0,1]$, construct some set of integers S_a . This means we have an uncountably large collection of sets of integers. Figure out a way to construct these sets S_a such that they satisfy the following property: if $a < b$, then S_a is a strict subset of S_b .

And that's all there is for this chapter on infinities.

Real Analysis

Chapter 3
Topology

Real Analysis

Chapter 3
Topology

Part 1
Metric Spaces



What counts as “distance”?

Some measures of distance are the “as the crow flies” distance function and the “manhattan” distance function. But what specifically makes a function $d(x,y)$ a valid way to measure distance between x and y ?



If someone claims that some number is the “distance” between A and B , It seems reasonable that we would want that distance to be the shortest way to get from A to B .

$$d(x,y) \leq d(x,z) + d(z,y)$$

(triangle inequality)



We also probably want distance to be same no matter where we started.

$$d(x,y) = d(y,x)$$

(symmetry)



And it would also be nice if distance was only 0 when you're measuring distance between something and itself.

$$d(x,y) = 0 \iff x=y$$


(identity of indiscernibles)

So, let's say a measure of “distance” is any function that satisfies all of the above criteria: triangle inequality, symmetry, and identity of indiscernibles.

Sets change a lot when you change the metric used to measure distances between them.



When measured with a different metric, the same set can seem like a different universe.



So, we often describe sets together with their metrics.

A metric space is any set accompanied by a distance metric used to measure the distance between its points.

For example, one really vanilla metric space is the set of real numbers measured by metric $d(x,y) = |x-y|$.

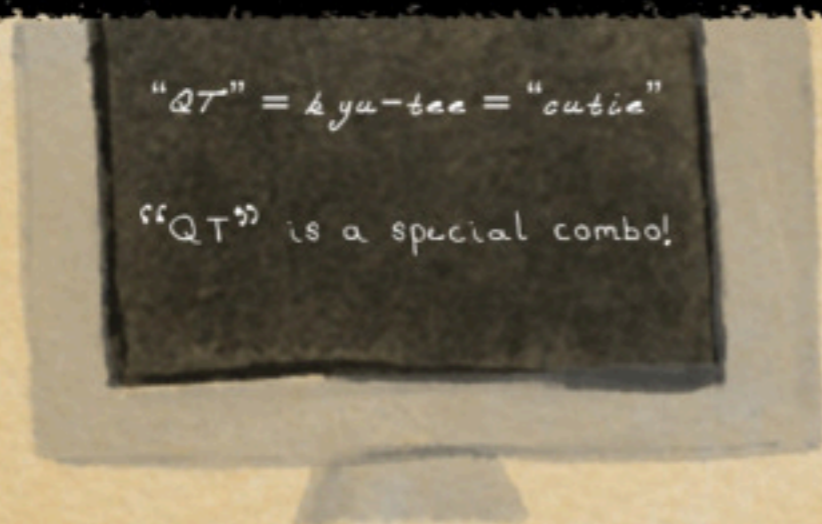
Here's a more interesting example of a metric space — a linguistic metric space. Say we wanted to find all words like QT (which sounds like "cutie") and OPM (which sounds like "opium"). That is, we want to find all letter combos that happen to sound like words.



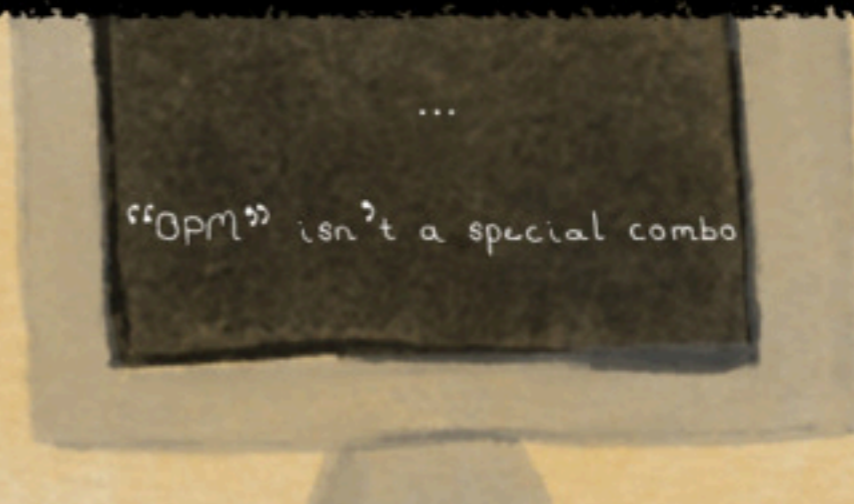
We could write a computer program to find all of them, by using a dictionary with a pronunciation of all the words, and another with a pronunciation of all the letters, and saving it to the list of "special combos" when we find a letter combination that matches a word in the dictionary.



This program does find "QT."



But it does not find "OPM."



Because OPM (oh-pee-ehm) doesn't sound exactly like opium (oh-pee-uhm).

To fix this, we need a way to say that the "em" sound is close enough to the "um" sound for us to accept it. One way to do this is by saying all like-sounding vowels are equal.




But just because vowel 1 sounds close enough to vowel 2, and vowel 2 sounds close enough to vowel 3, doesn't mean vowel 1 sounds close enough to vowel 3.



So vowel sounds aren't transitive. But equivalence is transitive (that is, if $a=b$, and $b=c$, then $a=c$). So we can't apply equivalence to vowel sounds.

However, distances, just like vowel sounds, aren't transitive...so what we really want is a distance metric.

...So what we really want is a distance metric. We can just create a set of vowel sounds...



And then we can use the "distance metric" that linguists have already determined to measure how far each one is from the other.

(If you're curious about how exactly linguists determine distances between vowels, I'd highly recommend the interactive Wikipedia page "IPA vowel chart with audio." And for a more in-depth look at the linguistics behind this, I'd recommend this detailed blog entry: <https://www.occasionalenthusiast.com/phonetic-distance-between-words-with-application-to-the-international-spelling-alphabet/>.)

Our new program can now identify words that sound "close enough" to a sequence of letters.



The entire point of this is: sets and their accompanying distance metrics exist outside regular-old 2d grids. That's one reason why it's useful to generalize the notion of a "distance" (like we did with our three-pronged definition a few pages back) — because now we can apply distances to more than just numbers.



Does teleportation violate notions of distance?

My favorite distance metric is one I like to call the "teleportation" metric.



In this metric, we imagine teleportation exists, and our current notion of distance does not. So, no matter where we want to travel, the place is 1 unit away. And if we want to stay in place, the distance is 0.


This metric can be written as:
$$d(x,y) = \begin{cases} 0 & \text{if } x = y \\ 1 & \text{otherwise} \end{cases}$$



Officially, this metric is called the "discrete" metric, because it involves no continuous spectrum of distances. The distance between everything is "discrete" — either a 0 or 1.

Brain teasers

1) Prove that the teleportation metric is an actual "metric" (a.k.a. "distance metric").




Hint:

Remember the properties that have to hold in order for a function to be a metric: the triangle inequality, symmetry, and identity of indiscernibles.

Answer:

I'll leave this up to you!

2) Come up with a function that isn't a valid distance metric (this is harder than it seems at first...).



Hint:

One such function is the function $d(x,y) = |x-y|$. Can you come up with another involving absolute values?

Answer:

One other such function is $|x-2y|$.

Real Analysis

Chapter 3
Topology

Part 2
Limit Points



This next chapter is about defining what a limit point actually is. From calculus class, already know what it is, sort of, in the context of functions.

But we hadn't really defined it rigorously. And rigor is what proving calculus is all about.

What's the difference between "really close" and "infinitely close"?



The finite set $\{.9\}$ gets really close, but not infinitely close to 1.

The finite set $\{.9, .99\}$ also doesn't get infinitely close.

Neither does the finite set $\{.9, .99, .999\}$.



...because we can always draw a neighborhood* around 1 that doesn't contain any of those points.

*Where we can formally define a "neighborhood" as the region within a distance of non-zero radius "r" around a center "c."

But the infinite set $\{.9, .99, .999\dots\}$ is infinitely close to 1. Because we can never draw a neighborhood around 1 that doesn't contain items in the set.



So let's say a **limit point** of a set (a point that the set gets infinitely close to) is a point such that every neighborhood around it contains a point of the set (besides the limit point itself*).

*Otherwise, every point would be a limit point, since every neighborhood around a point contains the point itself.


Can we make a set with exactly one limit point at 0?

So we want to draw a set...



...so that no matter what neighborhood we draw around 0, the neighborhood captures some points in the set.

One such set is $\{1/n: n \text{ is a natural number}\}$.



Another such set is $\{2/n: n \text{ is a natural number}\}$. And a bunch of others too.

Now, no matter what neighborhood you draw around "0", you end up with some items in the set.

Brain teaser

Can you make a set with exactly 3 limit points?



Your Thoughts Here:

Hint:

Try finding three different sets with exactly one limit point, and union-ing them.

Answer:

The set $\{1/n : n \in \mathbb{Z}^+\}$ (that is, the set of all numbers $1/n$, where n is a number in the set of positive integers) has exactly one limit point at 0. The set $\{2 + 1/n : n \in \mathbb{Z}^+\}$ has exactly one limit point at 1. The set $\{3 + 1/n : n \in \mathbb{Z}^+\}$ has exactly one limit point at 2. When we union those three sets together, we get a set with exactly three limit points.

Brain teaser

Draw a neighborhood around a limit point. Can that neighborhood ever have finitely many points in the set? Prove it.



Your Thoughts Here:

Hint:

Suppose there were a finite amount of points in the set. Can you draw a neighborhood that only contains the limit point?

But we know every neighborhood around the limit point must contain a point in the set. Contradiction.

Answer:

No. Each neighborhood around a limit point contains infinitely many points in the set. Proof: Say there were a finite number of points in some neighborhood around the limit point. Then you find the closest point to the limit point — say it's a distance " r " away. Now draw a neighborhood of radius slightly less than " r " around the limit point. This is a

Real Analysis

Chapter 3
Topology

Part 3
Open and Closed Sets



What are open sets?




Imagine a set where every single point can be perturbed a bit, in any direction, and still stay in the set.



Let's say such a set is "open."

*For example, the set $(0,1)$ — that is, the set of real numbers between 0 and 1, not inclusive — is an open set, because even a point as close to the edge as .01 can be moved a little bit (say, to .001) and still be in the set.

In more precise language, an open set is a set where each point has a neighborhood (of non-zero radius) around it that is contained entirely in the set. For example, the set $(-\infty, 0)$ is open. But the set $[0, 1)$ is not.



Note that in the real line, an "open set" is equivalent to an "open interval" such as (a,b) .

Is the union of open sets always open?

It turns out no matter how many open sets you combine together, the resulting set will always be an open set.

Proof: Say $S_1 \dots S_n$ are open sets. Then say the set $S = S_1 \cup S_2 \dots S_n$ is the union of all of them. Now, we consider a point p in the set S .




Consider the case $p \in S_1$.



Because S_1 is open, p has a neighborhood around it entirely within S_1 , and therefore entirely within this unioned set S .




Now consider the case $p \in S_2$.



Again, because S_2 is open, p has a neighborhood around it entirely within S_2 , and therefore entirely within S .



We can use a similar argument if p is an element of $S_3, S_4, \dots S_n$.




So, for any point p in this new set S , we can find a neighborhood around it entirely in the set S . So, S is open.

Is there anything interesting about the complement of open sets?

The complement of a set...

is all elements not in that set (but still in that metric space).



Do we notice any property that holds for every complement of every open set? We can look at examples for inspiration.

The complement of this open disk...



The complement of this open rectangle...



The complement of this set containing all reals < 0 ...




...is the boundary of that disk, and the surrounding area.



...is the boundary of that rectangle, and the surrounding area.




...is this set containing all reals ≥ 0 .



All the complements contain their limit points. Coincidence? Nope. It turns out, the complement of an open set always contains its limit points.

What is the complement of an open set?

So, since they are the complement of "open sets," let's say that any set that contains all its limit points is a "closed set."



Note that while the set $(0,1)$ is open because it doesn't contain its limit points 0 and 1, the set $[0,1]$ is closed.

Note that sometimes (but not always) these limit points will lie right on the boundary of the set, and so if the point is perturbed, it can fall right outside the set.



This is in contrast to open sets, where each point always had a bit of wiggle room (the required neighborhood around it).

Note that in the real line, an "open set" is equivalent to an "open interval" such as $[a,b]$.

“Hey, can you
check if my
backpack is
open?”

“You should
check if its
complement is
closed.”



How can you make an open set closed?



What's the smallest closed set that contains a given open set?




Some hints:

The set $(0,1)$ is open.

The set $[0,1]$ is closed.

The set of coordinates such that $0 < x < 1$ and $0 < y < 1$ is open.

The set of coordinates such that $0 \leq x \leq 1$ and $0 \leq y \leq 1$ is closed.



It turns out that given an open set S , the smallest closed set that contains it is the open set S plus the limit points of S .



Brain teaser

Is it possible for a set to be both open and closed? That is, can a set satisfy both:

(openness) every point has a neighborhood around it contained entirely within the set

(closedness) the set contains all its limit points



Your Thoughts Here:

Hint:

Is the set of all real numbers open? Is it closed?

Answer:

Yes. The set of all reals is both open and closed. So is the empty set (which has no points, so trivially, every point has a neighborhood around it within the set, and trivially, the set contains all its limit points). So maybe, "closed" and "open" weren't such great names. But what did we expect when we mixed math with something as complicated and contradictory as the English language?

Brain teaser

Is it possible for a set to be neither open nor closed? That is, can a set satisfy both:

(non-openness) the set has at least one point where you can't draw a neighborhood around it within the set

(non-closedness) the set doesn't contain at least one of its limit points

Your Thoughts Here:

Hint:

Consider a chunk of the real line that contains its limit point on one side, but not on the other.

Answer:

Yes. The set $(2, 3]$ is neither open nor closed. It's not open, because the point "3" is a point in the set that can't have a neighborhood around it entirely within the set. And it's not closed, because it doesn't contain one of its limit points: 2.

A bunch of open sets put together is called a "topology."

The rest of this chapter will be looking at properties of these "topologies." In particular, we'll be looking at the topological properties called compactness, separability, and connectedness.



Bibliography (Chapter 1)

For more on the origins of analysis as well as abstract mathematics in general, see the book "Duel at Dawn: Heroes, Martyrs, and the Rise of Modern Mathematics" by Amir Alexander.

Jerry Lembke's book on the Vietnam War spitting story is called "The Spitting Image: Myth, Memory and the Legacy of Vietnam." I owe this reference to Thalia Rubio at MIT.

For more on the Greek mathematician exiled for asserting that some numbers are irrational, see the Ted Ed video: "A Brief History of Banned Numbers."

For more on the contrast between the top-down historical development of analysis and the bottom-up teaching of analysis, see the book "How to Think about Analysis" by Lara Alcock.

I owe the question "what is a number" and its answer to "How We Got from There to Here: A Story of Real Analysis."

I owe this understanding of complex numbers to Prof. Francis Su's online lectures on Real Analysis.

Thanks to "How We Got from There to Here: A Story of Real Analysis" for an explanation of the Archimedean property of the reals.

I owe my understanding of how exponentials came to be defined to the books "Burn Math Class" by Jason Wilkes and "How We Got from There to Here: A Story of Real Analysis."

Many proofs and exercises are from "Principles of Mathematical Analysis" by Walter Rudin.

Bibliography (Chapter 2)

Thanks to Prof. Francis Su's online lectures on Real Analysis for explaining how our notion of bijections corresponds to our notion of counting.

For more on the development of attitudes towards infinity in mathematics, see "Infinity and the Mind: The Science and Philosophy of the Infinite" by Rudy Rucker.

Bibliography (Chapter 3)

Many proofs and exercises are from "Principles of Mathematical Analysis" by Walter Rudin.