Introduction

In mathematics, the simplest seeming ideas frequently draw forth entire worlds of complexity upon a closer look. Nothing seems to exemplify this more starkly than the continuum: a line that extends endlessly in both directions.

Often we distinguish a point in the "center" called the *origin* or *basepoint*. The complexity is revealed when we connect the continuum to the idea of number and the origin to zero.

Everything up to fractions is straightforward:

1. Distinguish a second point, let's say to the right of the origin, and connect it to the number one. There now exists a line segment on the line with endpoints zero and one. This segment is called the *unit interval*.

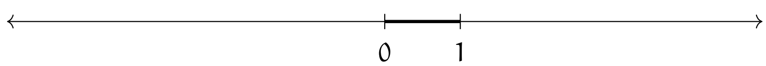

2. Generate the points connected to two, three, four, etc. by marching the unit interval rightwards along the line, so that the line segment from zero to one is the same length as the line segment from one to two, which is the same length as the line segment from two to three, and so on. There is no "largest" number, so these translated line segments cover the entire section of the line right of the origin, including the origin.
3. Generate the negative numbers by reflecting this right half of the line about the origin, so that there is a point connected to negative one exactly one unit interval's length to the left of the origin, a point connected to negative two exactly two unit interval lengths to the left of the origin, and so on.
4. The entire line is now covered with unit intervals, and the points where the unit intervals overlap (for example, two would be the overlap between the interval going from one to two and the interval going from two to three) form a uniform mesh.

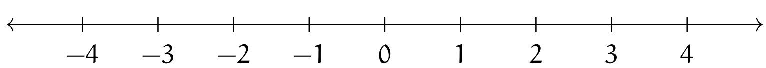

5. Shrink the mesh, so that it stays uniform, so that the origin stays put, and so that the point connected with the number x (which we will pick to the right of the one-point) moves to the point connected with one. We connect the point where the one-point ends up after this shrinking with the number $1/x$. This yields a line segment x times shorter than the unit interval, and we can march it to the right and left in order to generate a mesh x times finer than the unit mesh.

We have, at this point, assigned so many numbers to points on the line, in such an exhaustive way, that any point on the line is nearby (and can be brought nearer if needed) some point connected to a fraction.

The study of what happens in between the fraction-points is known as **real analysis**.

Existence of Irrational Numbers

Theorem 0.1. *If there is a number x with the property*

$$x \cdot x = 2,$$

then x is irrational.

Proof. Suppose we have such a number, and suppose further that that number can be written as a fraction p/q (written in lowest terms, so that p and q have no common factors). Then

$$2 = x \cdot x = (p/q) \cdot (p/q) = (p \cdot p)/(q \cdot q).$$

Multiply both sides by $q \cdot q$.

$$2 \cdot (q \cdot q) = p \cdot p \tag{1}$$

Then $p \cdot p$ is even. Suppose further that p is odd. Since the square of an odd number is odd, this cannot be the case: so p is even as well. Thus we may write $p = 2 \cdot r$ for some integer r. Substitute this into (1) to obtain

$$2 \cdot (q \cdot q) = (2 \cdot r) \cdot (2 \cdot r) = 4 \cdot (r \cdot r).$$

Cancel out a 2 on both sides, and we find that $q \cdot q$ is even. This is familiar: it's exactly what just happened with p. Following the logic,

$$q = 2 \cdot s$$

for some integer s. Note, however, that we have just witnessed two dividing both p and q. We assumed that p and q had no common factor, so the thing we supposed after the existence of x – that x can be written as a fraction – must be false. Either that, or x doesn't exist at all. ■

So, if we want numbers like x to exist, then we need a more extensive number system than the rationals.

Why should x exist? One answer is that we can search for it.

How To Find Square Roots

Both of the methods we present use the following fact about squaring positive rational numbers.

Theorem 0.2. *Let x and y be positive rational numbers. Then $x < y$ exactly when $x \cdot x < y \cdot y$.*

Proof. Suppose $x < y$. Then $c \cdot x < c \cdot y$ for every positive rational number c. Thus

$$x \cdot x < x \cdot y = y \cdot x < y \cdot y,$$

where $c = x$ in the first inequality and $c = y$ in the second inequality. This shows that $x < y$ implies $x \cdot x < y \cdot y$.

Also $x = y$ implies $x \cdot x = y \cdot y$, which when combined with the previous result yields $x \leq y$ implies $x \cdot x \leq y \cdot y$. Now suppose $x \geq y$. Then $y \leq x$, and by the second result, $y \cdot y \leq x \cdot x$, hence $x \cdot x \geq y \cdot y$. However, $x \geq y$ implies $x \cdot x \geq y \cdot y$ is the same statement as $x \cdot x < y \cdot y$ implies $x < y$, which completes the proof. ■

Thus, both squaring and taking square roots preserve the order of the positive rationals. In other words, if we have a lower bound $a \cdot a$ on $x \cdot x$ where a is rational, and an upper bound $b \cdot b$ on $x \cdot x$ where b is rational, then

$$a \cdot a < x \cdot x < b \cdot b$$

implies $a < x < b$. Stricter bounds on $x \cdot x$ yield stricter bounds on x.

The following pseudocode details a method, known as *binary search*, which finds the square root of 2.

```
binsearch(a_0, b_0, ε):
    set x_0 = 1/2(a_0 + b_0), x = x_0, (a_1, b_1) = x_0^2 > 2 ? (a_0, x_0) : (x_0, b_0);
    set x_1 = 1/2(a_1 + b_1), x = x_1;
    while(|x_i − x_{i−1}| ≥ ε):
        set (a_{i+1}, b_{i+1}) = x_0^2 > 2 ? (a_i, x_i) : (x_i, b_i);
        set x_{i+1} = 1/2(a_{i+1} + b_{i+1}), x = x_{i+1};
    return x;
```

The symbol-string "`X ? Y : Z`" should be read "`Y` if `X` and otherwise `Z`."

To get an idea of how fast this is, here are the values of a_i, b_i, x_i for `binsearch(1, 2, 0.01)`:

a_i	b_i	x_i	$\lvert x_i - x_{i-1}\rvert$
1	2	1.5	NULL
1	1.5	1.25	0.25
1.25	1.5	1.375	0.125
1.375	1.5	1.4375	0.0625
1.375	1.4375	1.40625	0.03125
1.40625	1.4375	1.4140625	0.0078125

(For reference: $\sqrt{2} \approx 1.41421356237$.)

Note that the term $|x_i - x_{i-1}|$ gets twice as small with each step. Further, as the index i grows larger, if $i < j$ then

$$\begin{aligned}|x_i - x_j| &= |x_i + (0 + \cdots + 0) - x_j| \\ &= |x_i + ((x_{i-1} - x_{i-1}) + \cdots + (x_{j+1} - x_{j+1})) - x_j| \\ &= |(x_i - x_{i-1}) + \cdots + (x_{j+1} - x_j)| \\ &\leq_! |x_i - x_{i-1}| + \cdots + |x_{j+1} - x_j| \\ &\leq (j - i) \cdot |x_i - x_{i-1}| \to 0.\end{aligned}$$

(The $\leq_!$ step uses a result called the *triangle inequality* – we explain further later on.)

Thus, the terms

$$x_0,\ x_1,\ x_2, \ldots$$

in the progression $(x_i)_i$ get closer and closer to each other as i gets large. This means that for very large i, not only is there not much difference between x_i and x_{i+1}, there also isn't much difference between x_i and $x_{i+10000000}$. Even simpler: far enough out, $(x_i)_i$ basically stops moving. This is stiking evidence that $(x_i)_i$ is actually zeroing in on some point on the line, even if that point does not associate with a rational number.

So, we at least have a candidate for $\sqrt{2}$.

Note, however, that we have not yet shown that potentially distinct methods produce the same result.

The binary search method of finding square roots works, but has its flaws: it is slow, and the iterates potentially dip below the target. Our next method will be faster, with the property that each iterate is strictly above the target. It even has simpler pseudocode. It was known to the Greeks (some say the Babylonians came up with it).

Hero's Method

We begin with a helpful fact:

Theorem 0.3 (AM–GM inequality). *Let a and b be nonnegative rational numbers. Then*

$$ab \leq \frac{(a+b)^2}{4}.$$

Proof. Observe that for every rational number x, the number x^2 is always either 0 or positive.

In particular, since $0 \leq (a-b)^2 = a^2 + b^2 - 2ab$, we get

$$4ab \leq a^2 + 2ab + b^2 = (a+b)^2$$

from which the claim follows. ■

The insight is that we can use x_i and $\frac{2}{x_i}$ for a and b above.

For example, if $x_{i+1} = \frac{1}{2}(x_i + \frac{2}{x_i})$, then this yields:

$$2 = x_i \cdot \frac{2}{x_i} \leq \frac{1}{4}\left(x_i + \frac{2}{x_i}\right)^2 = x_{i+1}^2$$

so that as long as the initial x_0 is positive, the squares of all subsequent terms will never exceed 2.

Here is the pseudocode for Hero's method:

```
hero(x_0, ε):
    set x_1 = 1/2(x_0 + 2/x_0), x = x_1;
    while (|x_i − x_{i−1}| ≥ ε):
        set x_{i+1} = 1/2(x_i + 2/x_i), x = x_{i+1};
    return x
```

Hero's method is fast (reminder: $\sqrt{2} \approx 1.41421356237$):

x_i	$\lvert x_i - x_{i-1} \rvert$
1	NULL
1.5	0.5
1.4167...	0.0833...
1.414215 69...	0.00245...
1.414213562...	0.000002...

A Comment on Speed

Binary search is slower than Hero's method, but still has this nice property that for i large, x_i and $x_{i+10000000}$ are about as close to each other as x_i and x_{i+1} are. So the fact that Hero's method is even faster means that the rate at which binary search and Hero's method match up is limited by the rate at which binary search reaches the square root of 2 at all.

We also note that Hero's method isn't even the fastest method known (cf. Householder's methods)!

How Many Line Numbers?

So, what happens if we start accepting all solutions to $P(\xi) = 0$ where $P(\xi) \in \mathbf{Q}[\xi]$ (here, ξ would be an *algebraic number*, i.e. a root of a polynomial with rational coefficients) as valid numbers? Well, the first thing to notice is that some of these solutions aren't even on the line anymore (consider $P(\xi) = \xi^2 + 1$). But even if we were to include these exterior values, we still wouldn't have enough numbers to fully exhaust the continuum.

Liouville Numbers

Definition 0.4. A **Liouville number** is a number λ such that

for any integer n, there exist integers p and q with $q > 1$ such that

$$0 < \left|\lambda - pq^{-1}\right| < q^{-n}.$$

An example of such a number would be

$$\lambda = \sum_{k \geq 1} \frac{a^k}{b^{k!}}.$$

Now, we're rapidly outpacing the number of answers we can supply rigorously with respect to the number of questions we generate. Here are some facts, though, to chew on:

- There are more Liouville numbers than there are rational or algebraic numbers in a precise sense: The set of Liouville numbers is uncountable, whereas the rational and algebraic numbers are countable.
- At the same time, "almost none" of the numbers on the line are Liouville numbers, i.e. the set of Liouville numbers has *measure zero*.
- All Liouville numbers are non-algebraic, or *transcendental*, by construction.

The very existence of Liouville numbers is a clue to just how vast the set of numbers on the number line may be.

The Continuum Hypothesis

Here's an even more disconcerting fact: philosophically, the jury is still out on how many numbers lie on the number line. This proposition, known as the **continuum hypothesis**, has been shown to be *independent* of the axioms upon which standard mathematics is built!

The Silver Lining

It seems that every time we reach an apex of crystallized understanding of the number line from its points, some random bit of chaos swoops in and shatters our progress. So, why are we confident that this number line object is worth studying at all?

The answer is that we can take the sheer uncountability of the number line and simply incorporate it into our toolkit. We can weaponize the very pathologies that afflicted us in ways that fortify our understanding instead of crippling it. Because while the real line may be algebraically incomprehensible, it just so happens to be topologically graspable.

So—forging bravely ahead—let $\mathbf{R}$ denote the set of numbers on the number line.

If a number x is in $\mathbf{R}$, we will call x a **real number**.