The point that refuses to move

By Zixuan Wang | 1997 words | 12 min read

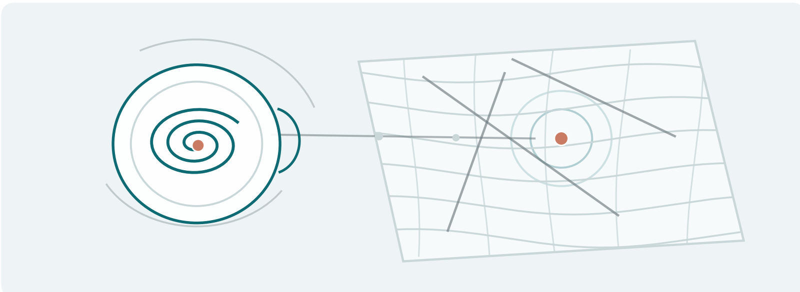

Introduction

If you stir a cup of coffee, must at least one drop of coffee necessarily end up at its original position? If you crumple up a map of England and lay it out randomly on the ground, will there be a point on the map that lies exactly above the real-world location it represents?

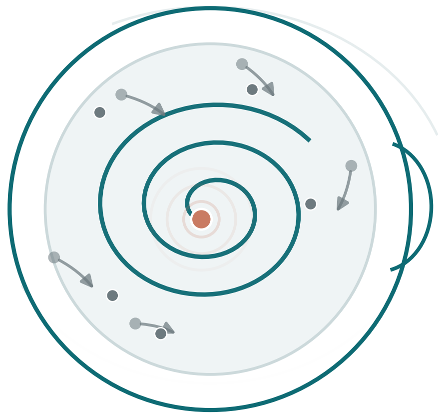

The Brouwer fixed-point theorem provides an answer that is counter-intuitive to many: yes. Roughly speaking, it asserts that any continuous way of moving points inside a closed, bounded shape must leave at least one point unmoved.

What is a fixed point?

A fixed point is simply a point that the function leaves where it is. That is all.

If f is a function from a set X to itself, then a *fixed point* of f is a point $x \in X$ such that

$$f(x) = x.$$

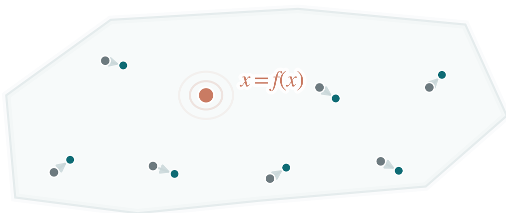

For example, if $f(x) = x^2$ on the real numbers, then 0 and 1 are fixed points, because $0^2 = 0$ and $1^2 = 1$. If $f(x) = x + 1$, there are no fixed points at all. Every point is shifted one step to the right; nothing stays put.

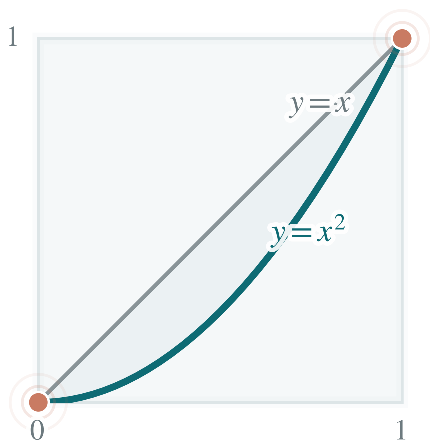

Brouwer’s theorem is not about arbitrary functions. It is about functions from a space back into itself which are continuous. In this setting, continuity means, informally, that nearby points are sent to nearby points. If you imagine the space as made of rubber, a continuous map deforms the rubber without tearing it or gluing distant parts together by fiat.

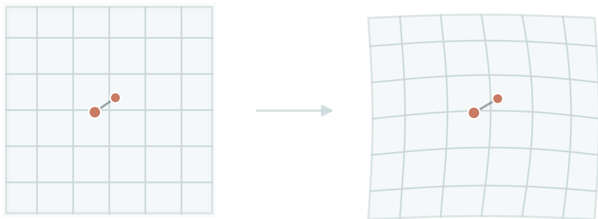

The one-dimensional case

The one-dimensional version of Brouwer’s theorem says:

If $f : [a, b] \to [a, b]$ is continuous, then there exists some $c \in [a, b]$ such that $f(c) = c$.

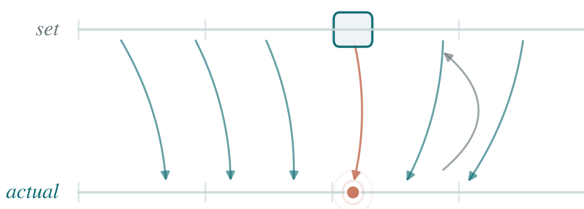

The key trick is to compare $f(x)$ with x itself. Define

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

Why do this? Because a fixed point of f is exactly a point where $g(x) = 0$. Indeed,

$$g(x) = 0 \iff f(x) - x = 0 \iff f(x) = x.$$

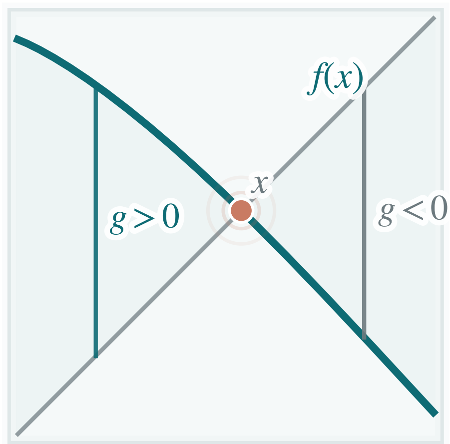

c

So the problem becomes: why must the continuous function g vanish somewhere?

Now use the fact that f maps $[a, b]$ into itself. Since $f(a) \in [a, b]$, we have $f(a) \geq a$. Therefore

$$g(a) = f(a) - a \geq 0.$$

Likewise, $f(b) \in [a, b]$, so $f(b) \leq b$, and hence

$$g(b) = f(b) - b \leq 0.$$

Thus g starts at the left endpoint with a value that is non-negative, and ends at the right endpoint with a value that is non-positive.

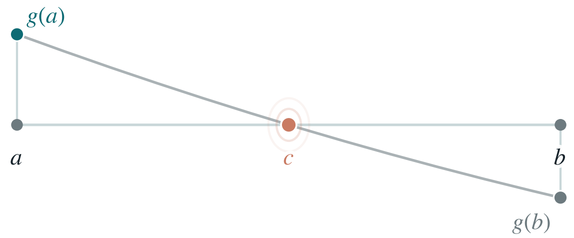

If one of these values is already 0, we are done: a or b is a fixed point. Otherwise $g(a) > 0$ and $g(b) < 0$. Since g is continuous, the Intermediate Value Theorem applies: a continuous function on an interval cannot move from positive to negative without passing through 0. So there exists $c \in (a, b)$ such that

$$g(c) = 0.$$

Equivalently,

$$f(c) = c.$$

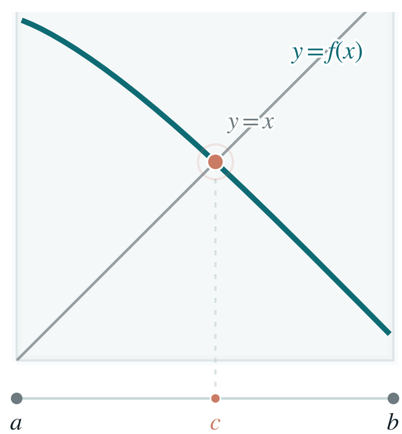

We do not compute the fixed point. We do not know where it is. We merely know it must exist, because the graph of $y = f(x)$ cannot stay entirely above or entirely below the diagonal line $y = x$ while remaining trapped inside the square $[a, b] \times [a, b]$. It would be lovely if higher dimensions succumbed as briskly.

Why the two-dimensional case is genuinely harder

For the closed disc

$$D^2 = \{(x, y) \in \mathbb{R}^2 : x^2 + y^2 \leq 1\},$$

Brouwer says: every continuous map $f : D^2 \to D^2$ has a fixed point.

And in general, for the closed unit ball

$$B^n = \{x \in \mathbb{R}^n : \|x\| \leq 1\},$$

the theorem says:

For every integer $n \geq 1$, every continuous map $f : B^n \to B^n$ has a fixed point.

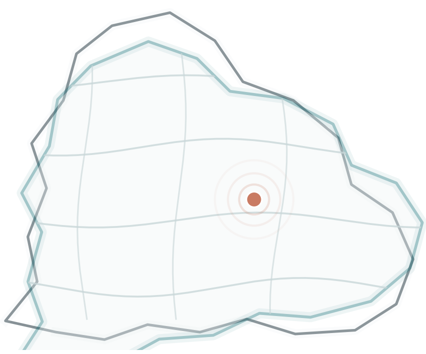

There is an equivalent version for the n-simplex, and that is the one we shall actually prove.

Why not repeat the one-dimensional trick? Why not set $g(x) = f(x) - x$ and look for $g(x) = 0$?

Because now $g(x)$ is not a real number but a vector in $\mathbb{R}^n$. In one dimension, saying a number is positive or negative gives you leverage; the Intermediate Value Theorem exploits the order of the real line. In higher dimensions there is no comparable order on vectors. A zero of the first coordinate of g may occur in one place, a zero of the second coordinate somewhere else. You need all coordinates to vanish at the same point, and there is no scalar trick of equal simplicity that forces this.

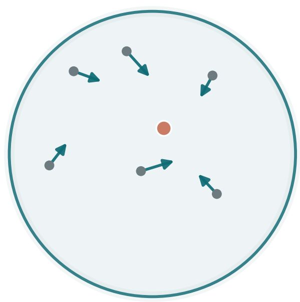

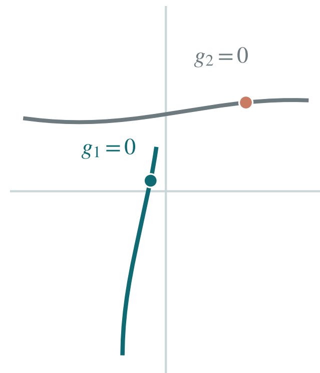

The hypotheses matter too. Remove closedness, and the theorem fails at once: the map

$$f : (0, 1) \to (0, 1), \qquad f(x) = \frac{x+1}{2}$$

has no fixed point, because solving $x = (x + 1)/2$ gives $x = 1$, which is not in the open interval. Remove boundedness, and it fails again: on $\mathbb{R}$,

$$f(x) = x + 1$$

has no fixed point. The same kind of thing happens in higher dimensions. Compactness — in Euclidean space, closed and bounded — is doing real work.

But compactness is not the whole story by itself. Not every closed bounded set has the fixed-point property. Rotate the circle S^1 by a non-zero angle and no point stays put. So Brouwer is not a theorem about arbitrary tidy sets. It is about ball-like, hole-free ones — and, more generally, about compact convex sets.

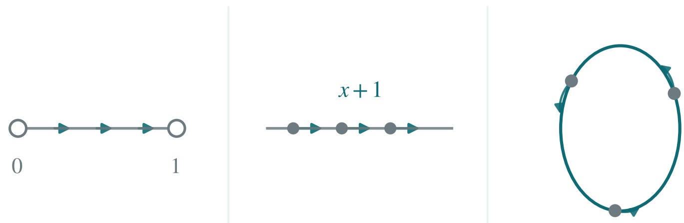

Why simplices are the right battleground

A 0-simplex is a point. A 1-simplex is an interval. A 2-simplex is a triangle. A 3-simplex is a tetrahedron. In general, the standard n-simplex is

$$\Delta^n = \left\{ (t_0, \dots, t_n) \in \mathbb{R}^{n+1} : t_i \geq 0 \text{ for all } i,\ \sum_{i=0}^{n} t_i = 1 \right\}.$$

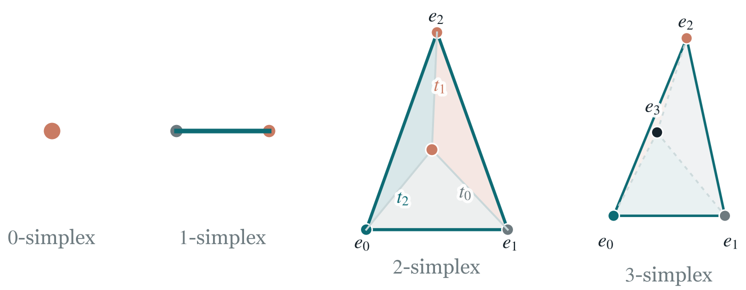

Its vertices are the standard basis vectors

$$e_0 = (1, 0, \dots, 0),\ e_1 = (0, 1, 0, \dots, 0),\ \dots,\ e_n = (0, \dots, 0, 1).$$

Every point $x \in \Delta^n$ can be written uniquely as

$$x = t_0 e_0 + \cdots + t_n e_n$$

with $t_i \geq 0$ and $t_0 + \cdots + t_n = 1$. These numbers $t_0, \dots, t_n$ are the barycentric coordinates of x. They tell you how much of each vertex is mixed into the point.

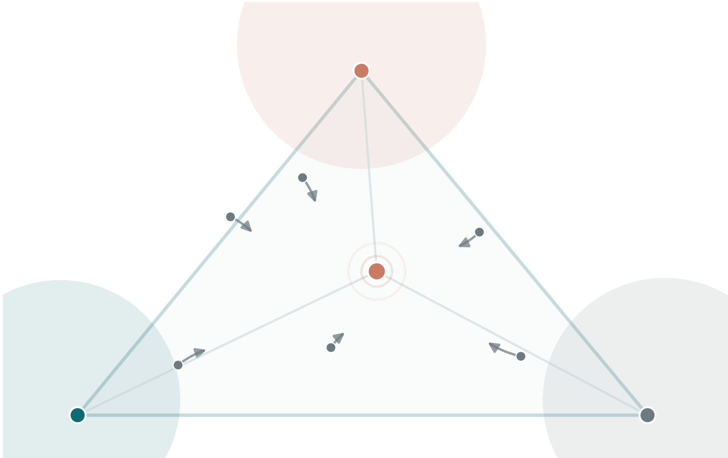

A face of Δ^n is what you get by declaring some of the barycentric coordinates to be 0. For example, the face spanned by vertices e_i with $i \in I \subseteq \{0, \dots, n\}$ consists of points with $t_j = 0$ whenever $j \notin I$.

Why prove Brouwer for Δ^n instead of B^n? Because simplices have flat faces and combinatorics. You can chop them into smaller simplices in a controlled way, and those little pieces can be labelled and counted.

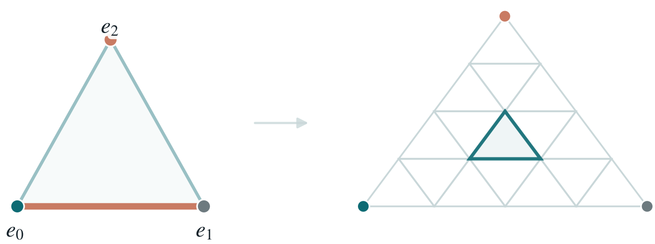

A triangulation of Δ^n is a finite decomposition of it into smaller n-simplices which fit together face-to-face: any two meet, if at all, along a common face. In a triangle this is just a dissection into little triangles; in a tetrahedron, into little tetrahedra. The mesh of a triangulation is the largest diameter of any small simplex. By refining the triangulation again and again, we can make the mesh tend to 0.

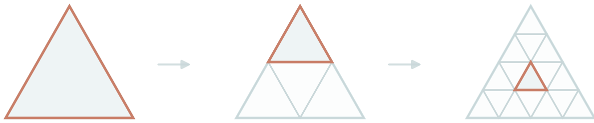

This is where the proof becomes unexpectedly combinatorial.

Sperner's Lemma

Suppose Δ^n has been triangulated. Label every vertex of the triangulation by one of the numbers $0, 1, \ldots, n$, but with a boundary rule: if a vertex lies on a face spanned by vertices indexed by some set I, then it may only receive a label from I.

In a triangle, that means a vertex on the edge between corners 0 and 1 may be labelled only 0 or 1, never 2. A corner itself must carry its own label. In higher dimensions it is the same rule: on any face, only the labels of that face are allowed.

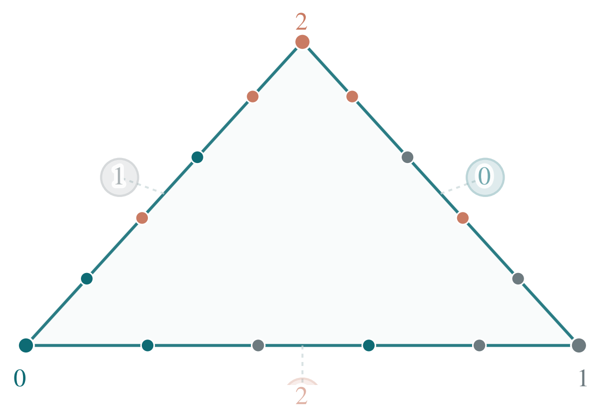

A small n-simplex in the triangulation is called fully labelled if its $n + 1$ vertices have all the labels $0, 1, \ldots, n$, one each.

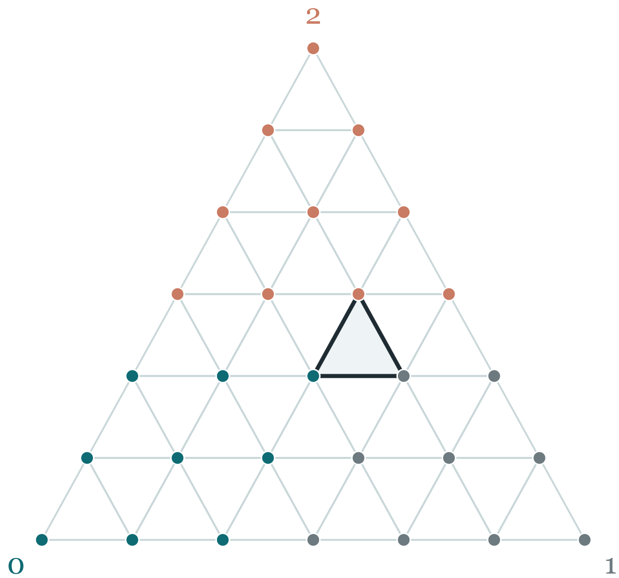

Now the remarkable claim.

Let T be a triangulation of Δ^n. Suppose each vertex v of T is assigned a label $L(v) \in \{0, 1, \dots, n\}$ such that whenever v lies on the face spanned by $\{e_i : i \in I\}$, its label lies in I. Then the number of fully labelled small n-simplices in T is odd. In particular, there is at least one.

When $n = 1$, the simplex is just an interval. Its two endpoints must be labelled 0 and 1. A triangulation is a subdivision into smaller intervals. As you walk from left to right through the sequence of vertices, the label starts at 0 and ends at 1. Each fully labelled small interval is exactly a place where the label changes between adjacent vertices. Since one starts at 0 and ends at 1, the number of changes is odd. So the lemma holds in dimension 1.

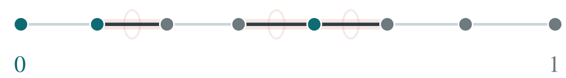

Now assume the lemma holds for dimension $n - 1$, and consider a Sperner-labelled triangulation T of Δ^n.

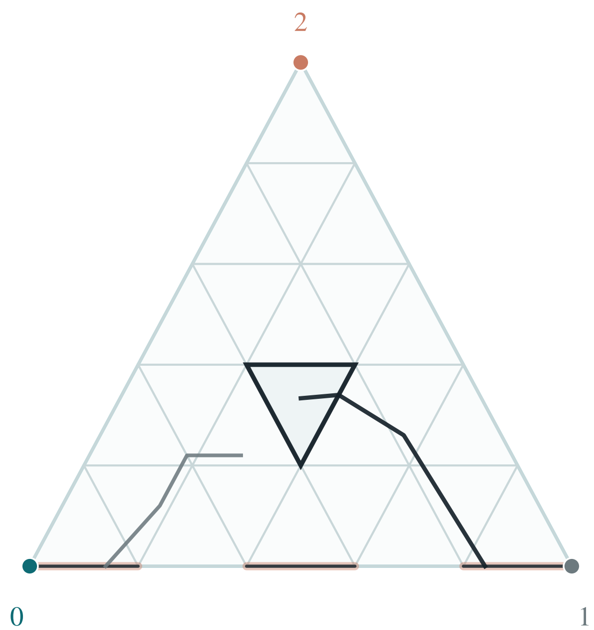

Focus on the face F opposite the vertex e_n. This is the face where the n-th barycentric coordinate is 0; it is itself an $(n - 1)$-simplex spanned by $e_0, \dots, e_{n-1}$. The triangulation T restricts to a triangulation of F, and the boundary condition on labels is still a Sperner condition there, now with labels only from $\{0, \dots, n - 1\}$. By the induction hypothesis, the number B of fully labelled $(n - 1)$-

simplices on F — meaning those whose labels are exactly $0, 1, \dots, n-1$ — is odd.

Now we perform a count in the whole triangulation, caring only about evenness or oddness.

Consider all $(n-1)$-faces of small n-simplices whose labels are exactly $0, 1, \dots, n-1$. Count each such face once for each small n-simplex containing it. Let the total count be I.

An interior $(n-1)$-face is shared by exactly two small n-simplices, so it contributes 2 to I, which is even and therefore irrelevant to parity. A boundary $(n-1)$-face is contained in exactly one small n-simplex, so it contributes 1. But a boundary face with labels $0, \dots, n-1$ can only lie on F: on any other boundary face one of those labels would be forbidden. Therefore the parity of I is exactly the parity of B. Since B is odd, I is odd.

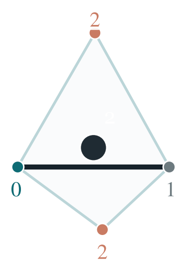

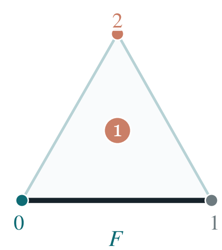

Now inspect how an individual small n-simplex can contribute to I.

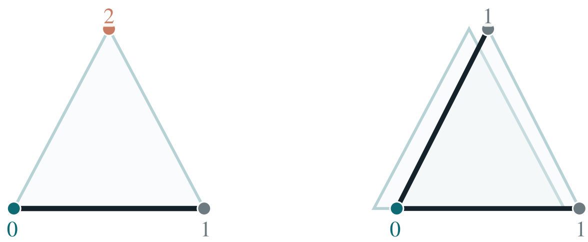

A small n-simplex has $n + 1$ vertices. For one of its facets — that is, one of its $(n - 1)$-dimensional faces — to have labels exactly $0, \dots, n - 1$, the labels on the $n + 1$ vertices of the whole simplex must be in one of two patterns.

Either the simplex is fully labelled, with labels $0, \dots, n$ each appearing once. Then exactly one facet has labels $0, \dots, n - 1$: the facet opposite the vertex labelled n.

Or the simplex uses only labels $0, \dots, n - 1$, with one of those labels repeated and label n absent. Then exactly two facets have labels $0, \dots, n - 1$: omit one or the other of the repeated vertices.

No other pattern contributes any such facet. So each small n-simplex contributes either 0, or 1, or 2 to I, and it contributes an odd number — namely 1 — exactly when it is fully labelled.

Hence the parity of I is the same as the parity of the number N of fully labelled small n-simplices. Since I is odd, N is odd. Therefore there is at least one fully labelled small simplex.

Constructing a Sperner Labelling

Now we use Sperner to prove Brouwer for the simplex.

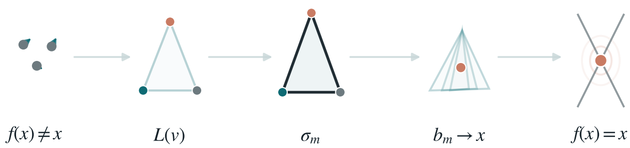

Assume, for contradiction, that there is a continuous map

$$f : \Delta^n \to \Delta^n$$

with no fixed point.

Write

$$f(x) = (f_0(x), f_1(x), \dots, f_n(x)).$$

Because $f(x) \in \Delta^n$, each coordinate $f_i(x)$ is non-negative and

$$f_0(x) + \cdots + f_n(x) = 1.$$

Take a sequence of triangulations $T_1, T_2, T_3, \dots$ of Δ^n whose mesh tends to 0. For each vertex v of each triangulation, we shall assign a label.

Since $f(v) \neq v$, the coordinates of $f(v)$ and v cannot all be equal. I claim that there must be at least one index i with

$$f_i(v) < v_i.$$

Why? Because if $f_i(v) \geq v_i$ for every i, then summing over all i gives

$$1 = f_0(v) + \cdots + f_n(v) \geq v_0 + \cdots + v_n = 1.$$

The inequality is actually an equality, so each coordinate must satisfy $f_i(v) = v_i$. That would mean $f(v) = v$, contradicting the assumption that there are no fixed points.

So at least one such index exists. Choose one of them and label the vertex v by that index i.

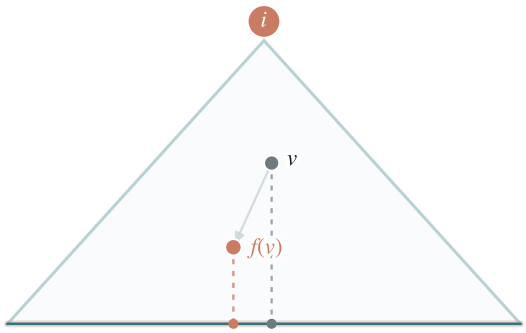

This labelling may not be unique, but it is well-defined in the only sense that matters: there is always at least one legal choice.

Now comes the boundary check. Suppose v lies on the face spanned by vertices indexed by a set I. That means $v_j = 0$ for every $j \notin I$. If $j \notin I$, then $f_j(v) \geq 0$, because $f(v)$ lies in the simplex. So it is impossible to have

$$f_j(v) < v_j = 0.$$

Hence no label outside I can ever be chosen. So the labelling obeys Sperner's boundary rule.

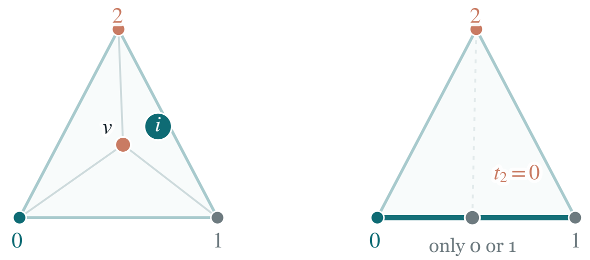

Therefore, by Sperner's Lemma, each triangulation T_m contains at least one fully labelled small n-simplex. Call one of them σ_m.

For each m, order the vertices of σ_m as

$$v_m^0, v_m^1, \dots, v_m^n$$

so that v_m^i has label i. Then by construction,

$$f_i(v_m^i) < (v_m^i)_i \qquad \text{for each } i = 0, \dots, n,$$

where $(v_m^i)_i$ means the i-th barycentric coordinate of the vertex v_m^i.

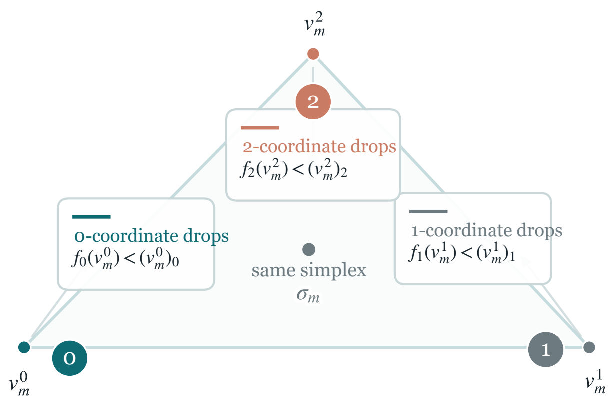

At first glance this is only approximate information. Coordinate 0 is smaller at one vertex, coordinate 1 at another, and so on. The crucial point of taking finer and finer triangulations is that these vertices all lie in a simplex whose diameter is going to 0. The separate witnesses are being squeezed together.

Compactness and the Final Contradiction

Let b_m be the barycentre of σ_m, the average of its vertices. Since Δ^n is closed and bounded in $\mathbb{R}^{n+1}$, it is compact. So the sequence (b_m) has a convergent subsequence. Pass to such a subsequence and write

$$b_m \to x \in \Delta^n.$$

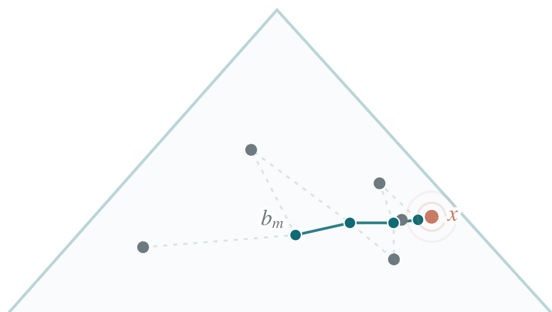

Because the diameter of σ_m is at most the mesh of T_m, and the mesh tends to 0, every vertex of σ_m lies closer and closer to its barycentre. Thus for each fixed i,

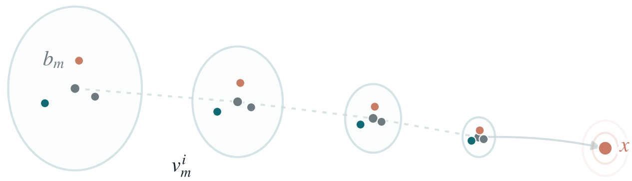

$$v_m^i \to x.$$

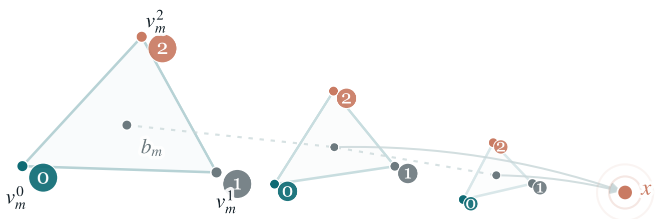

Now use continuity. From

$$f_i(v_m^i) < (v_m^i)_i$$

and the convergence $v_m^i \to x$, we obtain in the limit

$$f_i(x) \le x_i \qquad \text{for every } i = 0, \dots, n.$$

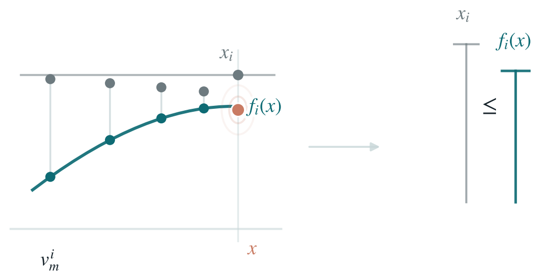

We have lost the strict inequality, which is no surprise: limits smooth out strictness. But the non-strict inequalities are enough. Summing them gives

$$f_0(x) + \cdots + f_n(x) \leq x_0 + \cdots + x_n.$$

Both sums are 1, because both $f(x)$ and x lie in the simplex. So in fact

$$1 \leq 1,$$

with equality throughout. If a collection of numbers satisfies $f_i(x) \leq x_i$ for every i and the total sums are equal, then each individual inequality must actually be an equality. Hence

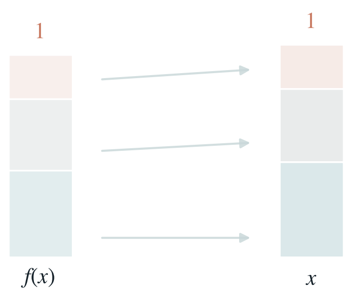

$$f_i(x) = x_i \qquad \text{for all } i.$$

Therefore

$$f(x) = x.$$

But that is a fixed point, contradicting our assumption that f had none.

So every continuous map $\Delta^n \to \Delta^n$ has a fixed point.

To pass from simplex to ball, choose a homeomorphism $h : \Delta^n \to B^n$: a continuous bijection with continuous inverse. Geometrically, you can obtain one by taking the barycentre of the simplex as a centre and sending each ray from that centre linearly onto the corresponding ray from the origin in the ball. If $F : B^n \to B^n$ is continuous, define

$$g = h^{-1} \circ F \circ h : \Delta^n \to \Delta^n.$$

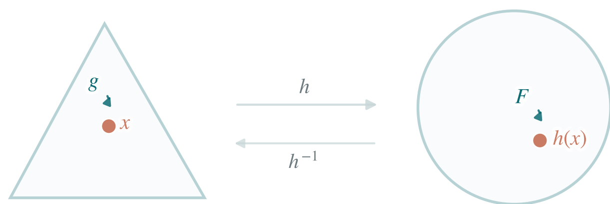

By the simplex case, g has a fixed point x. Then

$$h^{-1}(F(h(x))) = x,$$

so $F(h(x)) = h(x)$. Thus F has a fixed point in the ball. The disc case is simply $n = 2$, and the interval case $n = 1$.

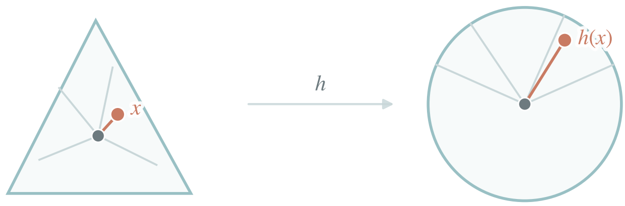

Conclusion

One could say, continuity may bend the world, but somewhere, something stubbornly stays put. Just like our passion for mathematics.