

Measuring the Unmeasurable

Hausdorff Dimension and Fractal Geometry

A Leisurely Excursion

Introduction

You probably learned in school that a line is one-dimensional, a square is two-dimensional, and a cube is three-dimensional. That is perfectly fine for the usual geometric objects you encounter every day. However, what do you do when the object in question is, well, weird?

Consider the Cantor set. You build it like this: start with the interval $[0, 1]$, remove the middle third, remove the middle third of each remaining piece, and repeat forever. What you're left with is a set that contains uncountably many points, yet has total length zero. It's not a finite collection of points (so it's not zero-dimensional in the usual sense), but it's too sparse to fill up any interval (so it's not really one-dimensional either). *So, what is it?*

This is the kind of question that Hausdorff dimension is designed to answer. The idea is that dimension doesn't have to be a whole number. A set can be $\log 2 / \log 3 \approx 0.63$ -dimensional, whatever that means. Our goal here is to make that precise. Along the way we'll need to rebuild our notion of "size" from scratch, so we'll first talk about outer measures and how to construct them, and then we'll get to the Hausdorff measure and dimension specifically.

Not Good Enough

Before going into the new machinery, let's ask: *why isn't the usual notion of length (or area, or volume) good enough?*

The problem is that ordinary n -dimensional Lebesgue measure is, in a sense, too rigid. Lebesgue measure is built to measure n -dimensional objects: it assigns length to curves, area to surfaces, volume to solids. But it has nothing useful to say about

objects that live between dimensions. If you try to measure the Cantor set with one-dimensional Lebesgue measure (length), you get zero. If you try to measure a curve in \mathbb{R}^2 with two-dimensional Lebesgue measure (area), you also get zero. The answer is always either zero or infinity, never something in between.

So the measure itself is not the issue. The issue is that we're using the wrong measure for the wrong object. The fix is to parametrize by dimension: instead of a single measure, we build a whole family of measures \mathcal{H}^s for $s \geq 0$, each designed to detect s -dimensional size. For an s -dimensional object, \mathcal{H}^s gives something finite and positive, while \mathcal{H}^t for $t \neq s$ gives either zero or infinity. This family is the Hausdorff measure.

Outer Measures

Let us take a small detour through the idea of an outer measure, since that's the language we'll use to set everything up.

Definition 1 (Outer Measure). An *outer measure* on a set X is a function $\mu^* : 2^X \rightarrow [0, \infty]$ satisfying:

1. $\mu^*(\emptyset) = 0$.
2. *Monotonicity*: if $A \subseteq B$, then $\mu^*(A) \leq \mu^*(B)$.
3. *Countable subadditivity*: for any countable collection $\{A_i\}$,

$$\mu^*\left(\bigcup_{i=1}^{\infty} A_i\right) \leq \sum_{i=1}^{\infty} \mu^*(A_i).$$

Notice that we don't require equality in the third condition, only \leq . An outer measure is defined on *all* subsets of X , not just on some restricted class of "measurable" sets. The price we pay for this generality is that the measure may not be additive on disjoint sets.¹

A natural way to build an outer measure is to cover sets with simpler pieces and measure those pieces. Specifically, if \mathcal{F} is a collection of sets (say, balls or cubes) and $\rho : \mathcal{F} \rightarrow [0, \infty]$ assigns a size to each, then

$$\mu^*(A) = \inf \left\{ \sum_{i=1}^{\infty} \rho(F_i) : A \subseteq \bigcup_{i=1}^{\infty} F_i, F_i \in \mathcal{F} \right\}$$

¹To recover additivity, we restrict to the class of *Carathéodory measurable* sets: A is measurable if $\mu^*(E) = \mu^*(E \cap A) + \mu^*(E \setminus A)$ for every $E \subseteq X$. One can show this class is a σ -algebra and μ^* is a genuine measure on it.

defines an outer measure on X . You cover A with pieces from \mathcal{F} , add up their sizes, and take the infimum² over all possible covers. This is pretty much how Lebesgue measure is defined, and it's also how Hausdorff measure is defined. The difference is in the choice of \mathcal{F} and ρ , and crucially, in a constraint on the *size* of the covering sets.

Hausdorff Measure

Here's the key idea. When we cover a set with balls of radius at most δ , the number of balls we need, and how their sizes contribute, changes as $\delta \rightarrow 0$ in a way that depends on the geometry of the set. We build this dependence into the definition by raising the diameter of each cover element to a power s , together with a normalization constant chosen so that the usual notions of length, area, and volume come out exactly right.

Definition 2 (Hausdorff Premeasure). Let $E \subseteq \mathbb{R}^n$, let $s \geq 0$, and let $\delta > 0$. Define

$$\mathcal{H}_\delta^s(E) = \inf \left\{ \sum_{i=1}^{\infty} \omega_s \left(\frac{\text{diam } U_i}{2} \right)^s : E \subseteq \bigcup_{i=1}^{\infty} U_i, \text{diam } U_i \leq \delta \right\},$$

where the infimum is over all countable covers of E by sets U_i with diameter³ at most δ , and

$$\omega_s = \frac{\pi^{s/2}}{\Gamma(\frac{s}{2} + 1)}$$

is the volume of the unit ball in “dimension” s .

Now, as δ decreases, we are forcing the covers to use smaller and smaller sets. This can only make the infimum larger: fewer covers are allowed, so we're taking the infimum over a smaller collection. Thus $\mathcal{H}_\delta^s(E)$ is non-decreasing in δ , and we can safely take a limit.

Definition 3 (Hausdorff Measure). The s -dimensional Hausdorff measure of E is

$$\mathcal{H}^s(E) = \lim_{\delta \rightarrow 0} \mathcal{H}_\delta^s(E) = \sup_{\delta > 0} \mathcal{H}_\delta^s(E).$$

One can verify that \mathcal{H}^s is indeed an outer measure on \mathbb{R}^n , and that Borel sets⁴ are Carathéodory measurable with respect to it. So we have a genuine measure on a very

²The *infimum* (written *inf*) of a set $S \subseteq \mathbb{R}$ is the greatest lower bound: the largest number that is \leq every element of S . Unlike the minimum, the infimum need not belong to S . For example, $\inf\{1/n : n \in \mathbb{N}\} = 0$, even though 0 is not in the set. In our context, $\inf\{\text{all cover costs}\}$ means the smallest total cost achievable, even if no single cover actually achieves it exactly.

³The *diameter* of a set $A \subseteq \mathbb{R}^n$ is $\text{diam}(A) = \sup\{|x - y| : x, y \in A\}$, the largest distance between any two of its points. For a ball of radius r the diameter is $2r$; for a unit square the diameter is $\sqrt{2}$.

⁴*Borel sets* form the smallest collection of subsets of \mathbb{R}^n that contains all open sets and is closed under countable unions and intersections. Every set you can write down explicitly (open sets, closed sets, their unions, intersections, and countable combinations thereof) is Borel. The collection is large enough to handle virtually all sets that arise in analysis.

large class of sets.

Let's check this against things we already know. When $s = n$ and $E \subseteq \mathbb{R}^n$, the n -dimensional Hausdorff measure \mathcal{H}^n agrees with n -dimensional Lebesgue measure.⁵ For $s = 1$, \mathcal{H}^1 measures the length of curves in \mathbb{R}^n , which is exactly what arc length does. For $s = 2$, \mathcal{H}^2 measures surface area. So the Hausdorff measure is a genuine generalization of the familiar notions, not some exotic replacement.

Before moving on, let's pause and think about what the two main ingredients of the definition are actually doing. There are two choices in that definition that might seem strange: why require $\text{diam}(U_i) \leq \delta$ at all, and why then take $\delta \rightarrow 0$? Both are essential.

Without a diameter constraint, you could cover any bounded set E with a single large ball of diameter D . That gives $\sum_i (\text{diam } U_i)^s = D^s$, which just depends on how large E 's bounding box is, not on the structure of E itself. A fractal set and a smooth curve sitting inside the same box would get the same bound. This is completely useless for us.

Forcing $\text{diam}(U_i) \leq \delta$ prevents this. The covering sets must now be small enough to actually follow the geometry of E . The more intricate E 's local structure, the more small sets you need, and the larger $\sum \omega_s (\text{diam } /2)^s$ becomes. See figure 1 to make sense of what we just talked about.

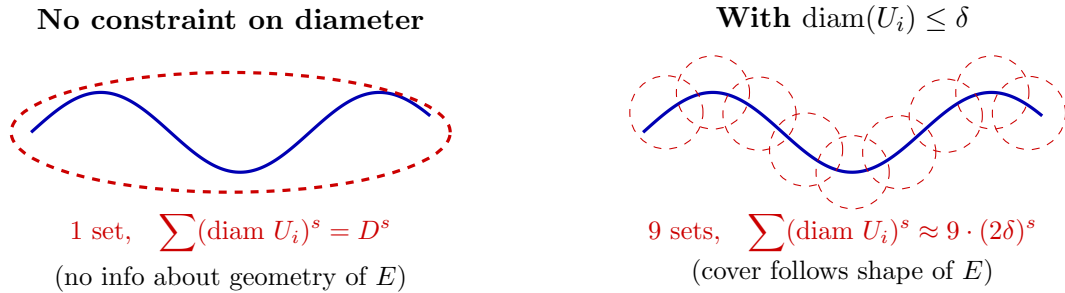


Figure 1: Without a diameter constraint, a single large ellipse covers E and the sum $\sum (\text{diam } U_i)^s$ carries no information about E 's geometry. With the constraint $\text{diam}(U_i) \leq \delta$, the covering sets are forced to hug the set, so the count and sizes of cover elements actually reflect the structure of E .

Now, why take $\delta \rightarrow 0$ at all? As δ decreases, fewer covers are admissible, so the infimum $\mathcal{H}_\delta^s(E)$ can only go up. It's non-decreasing in $\frac{1}{\delta}$. Taking the limit $\delta \rightarrow 0$ extracts the "true" s -dimensional size, rather than something that depends on the particular

⁵With the normalization above, one has $\mathcal{H}^n = \mathcal{L}^n$ on Borel subsets of \mathbb{R}^n . Many texts omit the factor $\omega_s 2^{-s}$ in the definition, in which case \mathcal{H}^n agrees with Lebesgue measure only up to a constant.

scale we happen to be looking at.

Here's the key calculation. If a d -dimensional set needs $\sim \delta^{-d}$ balls of radius δ to cover it, then

$$\sum_i (\text{diam } U_i)^s \sim \delta^{-d} \cdot \delta^s = \delta^{s-d}.$$

As $\delta \rightarrow 0$: this goes to 0 if $s > d$, blows up to ∞ if $s < d$, and stays finite if $s = d$. The limit $\delta \rightarrow 0$ is precisely what makes this whole thing sharp. Fix δ at any positive value, and you'll miss it.

The Critical Threshold

Here's where things get interesting. What happens to $\mathcal{H}^s(E)$ as s varies?

Theorem 1. *If $\mathcal{H}^s(E) < \infty$, then $\mathcal{H}^t(E) = 0$ for all $t > s$.*

Proof. Let $\{U_i\}$ be a δ -cover of E . For $t > s$,

$$\sum_i (\text{diam } U_i)^t = \sum_i (\text{diam } U_i)^{t-s} \cdot (\text{diam } U_i)^s \leq \delta^{t-s} \sum_i (\text{diam } U_i)^s.$$

Taking the infimum over covers and then letting $\delta \rightarrow 0$ gives $\mathcal{H}^t(E) \leq \lim_{\delta \rightarrow 0} \delta^{t-s} \mathcal{H}^s(E) = 0$. \square

So the function $s \mapsto \mathcal{H}^s(E)$ jumps from ∞ to 0 at some critical value. It can't oscillate, and it can't stay positive for two different values of s (unless one of them is the critical value itself). Figure 2 shows exactly this behavior. That transition point is exactly what we call the Hausdorff dimension.




Figure 2: As s increases, $\mathcal{H}^s(E)$ makes a single jump from ∞ to 0 at $s_0 = \dim_H(E)$. At the critical value itself, the measure may be 0, finite and positive, or ∞ .

Definition 4 (Hausdorff Dimension). The *Hausdorff dimension* of E is

$$\dim_H(E) = \inf\{s \geq 0 : \mathcal{H}^s(E) = 0\} = \sup\{s \geq 0 : \mathcal{H}^s(E) = \infty\}.$$

At the critical value $s = \dim_H(E)$ itself, the measure $\mathcal{H}^s(E)$ can be zero, positive and finite, or infinite. All three are possible, and figuring out which one holds for a specific set can be quite hard.

Remark 1. The Hausdorff dimension of a set in \mathbb{R}^n is always between 0 and n . A single point has Hausdorff dimension 0. A smooth k -dimensional submanifold of \mathbb{R}^n has Hausdorff dimension k . So for “nice” sets, Hausdorff dimension agrees with topological dimension. The interesting case is when they disagree.

It’s worth pausing to understand intuitively why covering by small balls detects dimension.

Think about covering a line segment of length 1 with balls of radius δ . You need roughly $\frac{1}{\delta}$ balls, each of diameter 2δ . So $\sum(\text{diam})^s \approx (1/\delta) \cdot (2\delta)^s = 2^s \delta^{s-1}$. As $\delta \rightarrow 0$, this goes to 0 if $s > 1$, to ∞ if $s < 1$, and to ∞ if $s = 1$. So the line segment has Hausdorff dimension exactly 1. Okay, this is good.

Now cover a square of side 1 with balls of radius δ . You need roughly $\frac{1}{\delta^2}$ balls. So $\sum(\text{diam})^s \approx (1/\delta^2) \cdot (2\delta)^s = 2^s \delta^{s-2}$. This goes to 0 if $s > 2$, and to ∞ if $s < 2$. So the square has Hausdorff dimension 2.

The pattern is clear: the dimension is exactly the exponent that compensates the polynomial growth of the covering number. For a “ d -dimensional” set, you need $\sim \delta^{-d}$ balls of radius δ , and the critical exponent is $s = d$.

The Cantor Set

Let’s compute the Hausdorff dimension of the middle-thirds Cantor set C .

Recall how C is built. Let $C_0 = [0, 1]$. To get C_1 , remove the middle third, leaving two intervals of length $1/3$. To get C_k , remove the middle third of each remaining interval. At stage k , we have 2^k intervals each of length 3^{-k} . The Cantor set is $C = \bigcap_{k=0}^{\infty} C_k$.

Figure 3 shows the first few stages of this construction alongside the natural δ_k -cover at each level.




Figure 3: The Cantor set construction as a sequence of δ_k -covers. At stage k , the natural cover uses 2^k intervals each of diameter 3^{-k} . As $k \rightarrow \infty$ (i.e., $\delta_k = 3^{-k} \rightarrow 0$), the covers get finer and the total cost $\sum_i (\text{diam } U_i)^s = (2/3^s)^k$ decays to 0 exactly when $s > \log 2 / \log 3$, pinpointing the Hausdorff dimension.

At each stage k , the collection of 2^k intervals of length 3^{-k} forms a 3^{-k} -cover of C . So for any $s > 0$,

$$\mathcal{H}_{3^{-k}}^s(C) \leq 2^k \cdot (3^{-k})^s = \left(\frac{2}{3^s}\right)^k.$$

If $s > \frac{\log 2}{\log 3}$, then $\frac{2}{3^s} < 1$, so this goes to 0 as $k \rightarrow \infty$, giving $\mathcal{H}^s(C) = 0$.

This tells us that $\dim_H(C) \leq \log 2 / \log 3$. To show the dimension is exactly $\log 2 / \log 3$, we need to show $\mathcal{H}^s(C) > 0$ for $s < \log 2 / \log 3$, which requires a lower bound on every cover. This is the harder direction.

Theorem 2. $\dim_H(C) = \frac{\log 2}{\log 3}$.

The key insight for the lower bound is that C supports a natural probability measure μ that is “spread evenly” across C in the following sense: for any interval I ,

$$\mu(I) \leq |I|^s, \quad s = \frac{\log 2}{\log 3},$$

where $|I|$ denotes the length of I .⁶ Once we have this, for any cover $\{U_i\}$ of C ,

$$1 = \mu(C) \leq \mu\left(\bigcup_i U_i\right) \leq \sum_i \mu(U_i) \leq \sum_i |U_i|^s,$$

so $\sum_i |U_i|^s \geq 1 > 0$ for every cover, which means $\mathcal{H}^s(C) \geq 1$.

So, $\mathcal{H}^s(C) \in [1, \infty)$ for $s = \log 2 / \log 3$, and $\dim_H(C) = \log 2 / \log 3 \approx 0.63$. The Cantor set really is between 0 and 1 dimensional.

⁶This measure is the so-called *Cantor measure* or *devil’s staircase measure*. It assigns measure 2^{-k} to each of the 2^k intervals in C_k , and these are consistent across levels.

Self-Similarity

A natural follow up is: *is there a faster way to compute Hausdorff dimension, without going through the measure directly?*

For self-similar sets like the Cantor set, the answer is yes. The Cantor set is the union of two scaled copies of itself, each scaled by a factor of $1/3$:

$$C = \frac{1}{3}C \cup \left(\frac{2}{3} + \frac{1}{3}C \right).$$

More generally, a self-similar set satisfying the *open set condition*⁷ is made up of N copies scaled by a common ratio r . For such sets, the Hausdorff dimension is the unique solution s to the equation

$$N \cdot r^s = 1, \quad \text{i.e.,} \quad s = \frac{\log N}{\log(\frac{1}{r})}.$$

For the Cantor set, $N = 2$ and $r = \frac{1}{3}$, giving $s = \frac{\log 2}{\log 3}$, confirming what we computed directly. For the Sierpiński triangle ($N = 3$, $r = \frac{1}{2}$), we get $s = \frac{\log 3}{\log 2} \approx 1.58$. It's more than one-dimensional but less than two-dimensional, which feels right since it fills part of the plane but has area zero.

Comparison with Topological Dimension


We should check that Hausdorff dimension at least agrees with what we already know for nice sets. For any set $E \subseteq \mathbb{R}^n$,

$$\dim_T(E) \leq \dim_H(E) \leq n,$$


where \dim_T denotes the topological dimension.⁸ So Hausdorff dimension never goes below topological dimension. A space that's one-dimensional in the topological sense can't have Hausdorff dimension less than 1. However, Hausdorff dimension can exceed topological dimension. The Cantor set has topological dimension 0 (it's totally disconnected, with no connected subsets other than points) but Hausdorff dimension $\frac{\log 2}{\log 3} > 0$. Space-filling curves have topological dimension 1 but Hausdorff dimension 2, because they genuinely trace out area even though they're the image of an interval.

⁷The open set condition says there exists a bounded open set U such that the scaled copies of U are disjoint subsets of U . For the Cantor set, take $U = (0, 1)$.


⁸Topological dimension is defined inductively: the empty set has dimension -1 , a space has dimension $\leq n$ if every point has arbitrarily small neighborhoods whose boundaries have dimension $\leq n - 1$. Points have dimension 0, curves have dimension 1, surfaces have dimension 2, etc.




(a) Stage 0



(b) Stage 1



(c) Stage 2



(d) Stage 3

Figure 4: First four iterations of the Sierpiński triangle construction. It is built from 3 copies of itself, each scaled by a factor of $1/2$, so its Hausdorff dimension is $\log 3 / \log 2 \approx 1.58$.

Another classical example is the Koch snowflake boundary. Topologically it is still a curve, so its topological dimension is 1, but its self-similarity gives

$$\dim_H(\partial K) = \frac{\log 4}{\log 3} \approx 1.26.$$

So it sits strictly between an ordinary curve and a genuine surface.




Figure 5: The Koch snowflake boundary after one refinement step. Each line segment is replaced by 4 segments of one-third the length, leading to Hausdorff dimension $\log 4 / \log 3 \approx 1.26$ while the topological dimension remains 1.

This gap between topological and Hausdorff dimension is, in many ways, the definition of a fractal. Objects where the two notions disagree are the ones where the usual geometric intuition breaks down. Hausdorff dimension is not determined by the topology of the space alone. You can find Cantor-like sets of any Hausdorff dimension in $[0, 1]$. For instance, just vary the ratio r at each step instead of always removing the middle third. All of them are homeomorphic⁹ to the standard Cantor set (they're all compact, metrizable, totally disconnected, perfect), yet their Hausdorff dimensions vary.

⁹Two spaces are *homeomorphic* if there is a continuous bijection between them whose inverse is also continuous (a rubber-sheet equivalence: stretching and bending are fine, tearing and gluing are not). Homeomorphic spaces share all topological properties, which is exactly why two homeomorphic Cantor-like sets can have different Hausdorff dimensions: Hausdorff dimension is not a topological invariant.