AN INTRODUCTION TO FRACTAL DIMENSION WITH APPLICATIONS TO BROWNIAN MOTION

GAVIN AN

ABSTRACT. Defining and analyzing the dimension of fractals is a subject that can lead to many unintuitive and elegant results. One particularly interesting application lies in a stochastic process known as Brownian motion.

Beyond simply defining fractal dimension and calculating it for various sets, this paper also studies other properties of Brownian motion and fractal dimension itself: the relationship between fractal dimension and self-similarity; the recurrence and transience of Brownian motion; and the zero set of Brownian motion, to name a few.

CONTENTS

1. INTRODUCTION

When first confronted with the idea of fractal dimension (i.e. a notion of dimension which assigns to some sets a non-integer value), one could be forgiven for being confused: how can something be "between" two and three dimensions? It all just seems rather contrived.

Of course, it is contrived—just like everything else in math. The question is whether mathematicians find the notion of fractal dimension useful in analyzing sets, and as it turns out, the answer is yes (for example in probability theory, where certain processes that draw inspiration from physics tend to display significant self-similarity). To better understand what is meant by fractal dimension, it may be helpful to think of dimension as representative of the complexity of a set. The Sierpiński triangle, for example, can be thought of as in some sense more complex than a one-dimensional curve (since its "coarseness" is infinite; zooming

Date: August 2025.
Supported by the UChicago Math REU program and mentored by Brin Harper.

in on the curve indefinitely will not render it smooth) but less complex than a two-dimensional shape (since, similarly to a line, it does not take up any area). Thus fractal dimension allows us to analyze and compare the complexity of different sets as well as the spaces in which they live.

In this paper, we rigorously define fractal dimension, deduce some properties of it, and apply it to analyze Brownian motion. In the following section, we go through some measure theoretical preliminaries. In Section 3, we construct Hausdorff dimension and calculate it for a few example sets. We follow Chapters 6 and 7 in Stein and Shakarchi's *Real Analysis* textbook here. [1] In Section 4, we construct Brownian motion (an important example displaying fractal qualities) and prove a few of its properties. In Section 5, we examine Brownian motion's zero set, and in Section 6, we explore some of the properties of multi-dimensional Brownian motion. In these sections, we particularly examine the fractal qualities of these mathematical objects, following Lawler's *Introduction to Stochastic Processes*. [2]

2. Preliminaries

We assume that the reader has a background in measure theory, such as the definition of the Lebesgue measure on $\mathbb{R}^d$. However, we will briefly review the ideas found in abstract measure theory, which we will use when defining fractal dimension. We loosely follow Chapter 6 of [1] here.

A **measure space** consists of a set X, a σ-algebra $\mathcal{M}$ of "measurable sets" in X, and a measure $\mu : \mathcal{M} \to [0, \infty]$ satisfying countable additivity. As such, it is often denoted by its parts $(X, \mathcal{M}, \mu)$. However, where there is no ambiguity, we will instead simply use X for brevity.

Let X be a set. A function μ_* on X that maps every subset of X to a value in $[0, \infty]$ is called an **exterior measure** if it satisfies the following properties:

(i) $\mu_*(\emptyset) = 0$.
(ii) (Monotonicity) If $E_1 \subset E_2$, then $\mu_*(E_1) \leq \mu_*(E_2)$.
(iii) (Countable subadditivity) If $E_1, E_2, \cdots$ is a countable family of sets, then

$$\mu_* \left(\bigcup_{j=1}^{\infty} E_j \right) \leq \sum_{j=1}^{\infty} \mu_*(E_j).$$

Given some exterior measure μ_* on X, we say a set $E \subset X$ is **Carathéodory measurable**, or simply **measurable**, if

$$\mu_*(A) = \mu_*(E \cap A) + \mu_*(E^c \cap A) \quad \text{for every } A \subset X.$$

Proposition 2.1. ([1], Theorem 1.1, Chapter 6) *Given an exterior measure μ_* on a set X, the collection $\mathcal{M}$ of Carathéodory measurable sets forms a σ-algebra. Moreover, μ_* restricted to $\mathcal{M}$ is a measure.*

We wish to make one further conclusion relating measurable sets to Borel sets (defined below) under certain circumstances. We say that a set X is a **metric space** if there exists a function $d : X \times X \to [0, \infty)$ such that

(i) $d(x, y) = 0$ if and only if $x = y$,
(ii) $d(x, y) = d(y, x)$ for all $x, y \in X$, and
(iii) $d(x, z) \leq d(x, y) + d(y, z)$ for all $x, y, z \in X$.

We say that such a function d is a **metric**. Using d, it is possible to define the open and closed sets of X. We can then define the Borel σ-algebra $\mathcal{B}_x$ as the smallest

σ-algebra (a collection of sets closed under countable unions, intersections, and complements) in X that contains the open sets of X. We say that a set in $\mathcal{B}_x$ is a **Borel set**. Now, we define the **distance between two sets** $A, B \subset X$ as

$$d(A, B) = \inf\{d(x, y) \mid x \in A, y \in B\}.$$

Then we say an exterior measure μ_* is a **metric exterior measure** if

$$\mu_*(A \cup B) = \mu_*(A) + \mu_*(B) \quad \text{whenever } d(A, B) > 0.$$

Proposition 2.2. ([1], Theorem 1.2, Chapter 6) *If μ_* is a metric exterior measure on a metric space X, then μ_* restricted to the Borel sets is a measure.*

As an example, $\mathbb{R}^d$ with the Lebesgue measure defined on the Borel sets is a measure space.

3. Hausdorff dimension

Before defining Hausdorff dimension, we must first define a new measure on $\mathbb{R}^d$ deeply related to the idea of dimension. We follow Chapter 7 of [1] in this section. We define the **exterior α-dimensional Hausdorff measure** of a set E as

$$m_\alpha^*(E) = \lim_{\delta \to \infty} \inf\{\sum_{k=1}^{\infty} (\text{diam } F_k)^\alpha \mid E \subset \bigcup_{k=1}^{\infty} F_k, \text{diam} F_k \leq \delta\},$$

where $\{F_k\}$ is a countable covering of sets whose **diameter**, defined as

$$\text{diam } F_k = \sup\{|x - y| \mid x, y \in F_k\},$$

is less than or equal to δ. It should be noted that

$$\mathcal{H}_\alpha^\delta(E) = \inf\{\sum_{k=1}^{\infty} (\text{diam } F_k)^\alpha \mid \bigcup_{k=1}^{\infty} F_k \supset E, \text{diam} F_k \leq \delta\}$$

increases monotonically as δ decreases, so the limit

$$m_\alpha^*(E) = \lim_{\delta \to \infty} \mathcal{H}_\alpha(E)$$

exists, though could be infinite.

Exterior Hausdorff measure can be shown to be a metric exterior measure, as defined in the previous section. Therefore, we can restrict it to the Borel sets and attain a measure called **Hausdorff measure**. Thus, we have created a measure space $(\mathbb{R}^d, \mathcal{B}_x, \mu_\alpha)$. Hausdorff measure satisfies the following property.

Proposition 3.1. ([1], Property 8, Chapter 7) *If $m_\alpha(E) < \infty$ and $\beta > \alpha$, then $m_\beta(E) = 0$. Also, if $m_\alpha(E) > 0$ and $\alpha > \beta$, then $m_\beta(E) = \infty$.*

Thus, for every set E, there exists some unique α such that

$$m_\beta(E) = \begin{cases} 0 & \beta < \alpha, \\ \infty & \beta > \alpha. \end{cases}$$

We define the **Hausdorff dimension** of E to be α. A set with a non-whole number dimension is called a **fractal**.

Now, we will go through some examples. The **Cantor set** is a fractal that can be constructed by recursively removing the middle-third intervals in $[0, 1]$. More specifically, if $\mathcal{C}_0 = [0, 1]$, we let

$$\mathcal{C}_1 = \left[0, \frac{1}{3}\right] \cup \left[\frac{2}{3}, 1\right], \quad \mathcal{C}_2 = \left[0, \frac{1}{9}\right] \cup \left[\frac{2}{9}, \frac{3}{9}\right] \cup \left[\frac{6}{9}, \frac{7}{9}\right] \cup \left[\frac{8}{9}, 1\right], \cdots$$

Figure 1. Seven iterations of the Cantor set[8]

where $\mathcal{C}_n$ is defined, given C_{n-1}, by removing the open middle third of each interval making up $\mathcal{C}_{n-1}$. The first seven iterations of the construction can be seen in Figure 1. We define the **Cantor set** $\mathcal{C}$ as

$$\mathcal{C} = \bigcap_{j=1}^{\infty} \mathcal{C}_j.$$

Instead of directly calculating the Hausdorff dimension of $\mathcal{C}$, we will prove a more general theorem which will allow us to deduce the dimension of many self-similar sets. This first requires the notion of self-similarity.

A function $S : \mathbb{R}^d \to \mathbb{R}^d$ is said to be a **similarity** with **ratio** $r > 0$ if

$$|S(x) - S(y)| = r|x - y|$$

for all $x, y \in \mathbb{R}^d$.

A set $F \subset \mathbb{R}^d$ is said to be **self-similar** if there exist finitely many similarities $S_1, S_2, \cdots, S_m$ with the same ratio r such that

$$F = S_1(F) \cup S_2(F) \cup \cdots \cup S_m(F).$$

Furthermore, we say that $S_1, S_2, \cdots, S_m$ are **separated** if there exists a bounded open set $\mathcal{O}$ such that

$$\mathcal{O} \supset S_1(\mathcal{O}) \cup S_2(\mathcal{O}) \cup \cdots \cup S_m(\mathcal{O})$$

and the $S_k(\mathcal{O})$ are disjoint. The set $\mathcal{O}$ is not necessarily related to F.

The Cantor set is self-similar in $\mathbb{R}$ with ratio $1/3$, since given similarities

$$S_1 = x/3, \quad S_2 = x/3 + 2/3,$$

we have

$$\mathcal{C} = S_1(\mathcal{C}) \cup S_2(\mathcal{C}).$$

Furthermore, with $\mathcal{O} = (0, 1)$, we can see that S_1 and S_2 are separated similarities. We will now state a few requisite lemmas for our theorem (with proofs found in the appendix). Throughout this section, we suppose $S_1, S_2, \cdots, S_m$ are m separated similarities with some contracting ratio r; that is, a ratio that satisfies $0 < r < 1$.

Lemma 3.2. *There exists a closed ball B such that $S_k(B) \subset B$ for all $k = 1, \cdots, m$.*

We will introduce a few pieces of notation for the next lemmas. Let $\tilde{S}(B)$ denote the set

$$\tilde{S}(B) = S_1(B) \cup S_2(B) \cup \cdots \cup S_m(B),$$

and $F_k = \tilde{S}^k(B)$ denote the set

$$F_k = \tilde{S}^k(B) = \tilde{S} \circ \cdots \circ \tilde{S}(B),$$

where k is the number of compositions. Also, for each $\delta > 0$ and set A, we let

$$A^\delta = \{x \mid d(x, A) < \delta\}.$$

For two compact sets A and B, we define the **Hausdorff distance** as

$$\operatorname{dist}(A,B) = \inf\{\delta \mid B \subset A^\delta \text{ and } A \subset B^\delta\}.$$

Lemma 3.3. *The distance function* dist *on compact sets in* $\mathbb{R}^d$ *satisfies the following properties:*

(i) $\operatorname{dist}(A,B) = 0$ *if and only if* $A = B$.
(ii) $\operatorname{dist}(A,B) = \operatorname{dist}(B,A)$.
(iii) $\operatorname{dist}(A,B) \le \operatorname{dist}(A,C) + \operatorname{dist}(C,B)$.
(iv) $\operatorname{dist}(\tilde{S}(A), \tilde{S}(B)) \le r\operatorname{dist}(A,B)$.

Lemma 3.4. *There exists a unique non-empty compact set F such that*

$$F = S_1(F) \cup \cdots \cup S_m(F).$$

The proof of this lemma, again found in the appendix, implies that every self-similar set F can be expressed uniquely as a countable intersection of unions of closed balls.

Theorem 3.5. *The Hausdorff dimension of the unique F defined in the previous lemma is* $\log m/\log(1/r)$.

Proof. Let $\alpha = \log m/\log(1/r)$. We will first prove that $m_\alpha(F) < \infty$, and then (in a more difficult step) prove that $m_\alpha(F) > 0$. This will prove that α is indeed the Hausdorff dimension of F by Proposition 3.1.

Recall that by Lemma 3.2, $F_k = \tilde{S}^k(B)$ is by construction the union of m^k sets with diameter cr^k (where $c = \operatorname{diam} B$), and that by our proof of that lemma (see appendix), each F_k is a covering of F. For notation sake, we will denote each of the m^k sets as

$$B_i, \quad 1 \le i \le m^k.$$

Then for any $\delta > 0$, we can select a sufficiently large k such that $cr^k < \delta$ to find

$$\mathcal{H}_\alpha^\delta \le \sum_{i=1}^{m^k} (\operatorname{diam} B_i)^\alpha \le c' m^k r^{\alpha k} = c', \tag{3.6}$$

since $mr^\alpha = 1$, because $\alpha = \log m/\log(1/r)$. Since c' is a constant that depends only on c, (3.6) holds for arbitrary δ, and as a result we have $m_\alpha(F) \le c'$.

Now, we prove $m_\alpha(F) > 0$. First, fix an arbitrary point $\overline{x}$ in F. We define the "vertices of the k^{th} generation" as the m^k points that lie in F and are given by

$$S_{n_1} \circ \cdots \circ S_{n_k}(\overline{x}), \quad \text{where } 1 \le n_1 \le m, \cdots, 1 \le n_k \le m.$$

We index each vertex by the similarities that resulted in it as $(n_1, \cdots, n_k)$. Note that vertices are not necessarily distinct. Then, fix $\mathcal{O}$ that satisfies the separation condition for the similarities $S_1, \cdots, S_m$. We define the "open sets of the k^{th} generation" to be the m^k sets given by

$$S_{n_1} \circ \cdots \circ S_{n_k}(\mathcal{O}), \quad \text{where } 1 \le n_1 \le m, \cdots, 1 \le n_k \le m.$$

The open sets of generation k are indexed by their similarities as $(n_1, \cdots, n_k)$. The open sets of the k^{th} generation are disjoint, since those of the first generation are. Consequently, if $k \ge l$, then each open set of the l^{th} generation contains m^{k-l} open sets of the k^{th} generation.

Let v be a vertex of the k^{th} generation, and let $\mathcal{O}(v)$ be the corresponding open set of the k^{th} generation such that both carry the same label $(n_1, n_2, \cdots, n_k)$. Since

$\overline{x}$ is at a fixed distance from the original set $\mathcal{O}$, whose diameter is finite and non-zero, we find that

(i) $d(v, \mathcal{O}(v)) \leq cr^k$,
(ii) $c'r^k \leq \text{diam}\mathcal{O}(v) \leq c''r^k$

for some constants $c, c', c'' > 0$. Our goal is to prove that for any $\delta > 0$ and countable covering $\cup_{j=1}^{\infty} F_j \supset F$ with $\text{diam}F_j < \delta$, we have

$$\sum_{j=1}^{\infty} (\text{diam}F_j)^{\alpha} \geq c > 0.$$

To do this, we first note that each F_j is contained in an open ball B_j of diameter twice $\text{diam}F_j$. Thus, we will scale the diameters of the balls by 2 and create a new open covering $\cup_{j=1}^{\infty} B_j \supset F$. Since F is compact, we can extract a finite subcover $\cup_{k=1}^{N} B_k \supset F$. Thus, we want to show that any such finite covering of balls whose diameters are less than δ satisfies

$$\sum_{k=1}^{N} (\text{diam}B_k)^{\alpha} \geq c > 0 \tag{3.7}$$

where c is a constant independent from δ and the cover of F chosen. If we can do this, then we will have shown that the analogous sum for finite covers taken from the original countable covering (whose diameters are less than δ) is also bounded below by a constant, and therefore that $m_{\alpha}(F) > 0$. Then suppose $\mathcal{B}$ is such a covering by balls; choose k such that

$$r^k \leq \min_{1 \leq j \leq N} \text{diam}B_j < r^{k-1}.$$

We will need one final lemma to prove our theorem, whose proof is again found in the appendix.

Lemma 3.8. *Suppose B is a ball in the covering $\mathcal{B}$ that satisfies*

$$r^l \leq \text{diam}B \leq r^{l-1} \quad \textit{for some } l \leq k.$$

Then B contains at most cm^{k-l} vertices of the k^{th} generation.

Now, let N_l denote the number of balls B_j in $\mathcal{B}$ such that

$$r^l \leq \text{diam}B_j \leq r^{l-1}.$$

By the lemma, we see that the total number of vertices of the k^{th} generation covered by $\mathcal{B}$ can be no more than $c\sum_{l=1}^{k} N_l m^{k-l}$. Since all m^k vertices of the k^{th} generation are contained in F, we must have $c\sum_l N_l m^{k-l} \geq m^k$, and hence

$$\sum_{l=1}^{k} N_l m^{-l} \geq 1/c.$$

Since $\alpha = \log m / \log(1/r)$, we see that $m^{-l} = r^{l\alpha}$, and so

$$\sum_{j=1}^{N} (\text{diam}B_j)^{\alpha} \geq \sum_{l=1}^{k} N_l r^{l\alpha} \geq 1/c,$$

because for each l, N_l balls have diameter at least r^l. Thus, we have (3.7) for all relevant cases, and $m_{\alpha}(F) > 0$. □

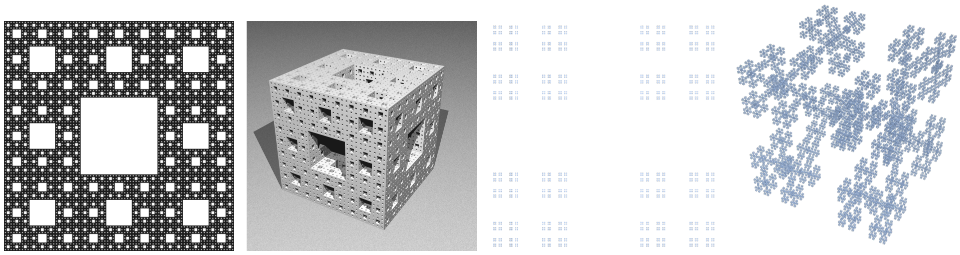

(A) Sierpiński carpet, iteration 6 (B) Menger sponge, iteration 4 (C) 2D Cantor dust, iteration 6 (D) 3D Cantor dust, iteration 5

FIGURE 2. Four fractals [8]

We can immediately apply this theorem to the Cantor set. Since $\mathcal{C}$ has $r = 1/3$ and $m = 2$, we have that $\dim \mathcal{C} = \log 2/\log 3$, or roughly 0.63. With a little extra effort, it can also be shown that the Cantor set has 0.63-dimensional Hausdorff measure 1. We will now look at three additional examples, as seen in Figure 2.

The first iteration of the Sierpiński carpet involves taking a square, dividing it into nine sub-squares, and removing the center one to create a ring. Each subsequent iteration repeats the process on the eight outer squares. [5] Thus, $m = 8$ and $r = 1/3$, so the dimension of the Sierpiński carpet is $\log 8/\log 3$, or roughly 1.89.

The Menger sponge takes this idea into three-dimensions by splitting a cube into twenty-seven sub-cubes and iteratively removing the center cube as well as the center of each of the six sides. [6] Thus, $m = 20$ and $r = 1/3$, so its dimension is $\log 20/\log 3$, about 2.73.

Cantor dust can be constructed in two-dimensions or three-dimensions. In two-dimensions, one divides a square as in the Sierpiński carpet, but removes all but the corner squares iteratively. In three-dimensions, a cube is divided as in the Menger sponge, but all but the corner cubes are removed iteratively. [7] The dimension of Cantor dust in two and three dimensions is therefore $\log 4/\log 3 \approx 1.26$ and $\log 8/\log 3 \approx 1.89$ respectively, since $r = 1/3$ and m is 3 or 8.

4. Brownian motion

A particularly interesting example of using Hausdorff dimension to analyze sets comes in Brownian motion, a type of stochastic process. More specifically, Brownian motion is a random continuous process. We will follow Sections 8.1 and 8.2 of [2] in defining and constructing this process.

Let X_t represent the value of a stochastic process at time t. We say that X_t is **Brownian motion** with mean 0 and variance σ^2 if:

(i) $X_0 = 0$,
(ii) For any $s_1 \leq t_1 \leq s_2 \leq t_2 \leq \cdots \leq s_n \leq t_n$, the random variables $X_{t_1} - X_{s_1}, \cdots, X_{t_n} - X_{s_n}$ are independent,
(iii) For any $s < t$, the random variable $X_t - X_s$ has a normal distribution with mean 0 and variance $(t-s)\sigma^2$, noting that **standard Brownian motion** has $\sigma^2 = 1$,
(iv) X_t is a continuous function of t; more formally, the function $t \to X_t$ is a continuous function.

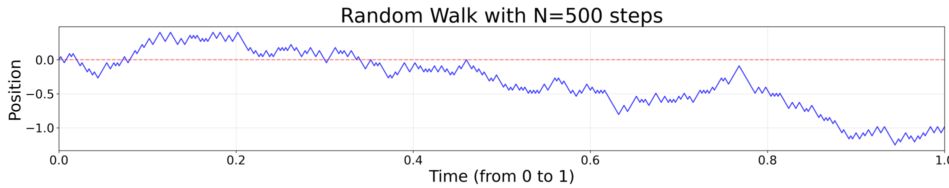

FIGURE 3. Random walk with $N = 500, a_N = 1/\sqrt{500}$ [9]

We can also refer to Brownian motion starting at x. If X_t is a Brownian motion as defined above, the motion Y_t constructed as

$$Y_t = X_t + x$$

is a Brownian motion starting at x.

This is a good definition of what Brownian motion should look like: it satisfies continuity, randomness, and has the correct mean and variance. Now, we will actually construct a process that satisfies it. There are several strategies for doing so, but we will take a limit of random walks.

A **random walk** can be thought of as the path of a person randomly walking down a number line. Letting S_t be our position at time t, we start at $S_0 = 0$. For each discrete increment in time (we will use 1 for now), we flip a coin. Heads, we increase our position by 1; tails, we decrease it by 1. Note that we are constrained to a single axis of motion. More formally, we have

$$S_n = Y_1 + Y_2 + \cdots Y_n,$$

where the Y_i are independent random variables such that

$$\mathbb{P}\{Y = 1\} = \mathbb{P}\{Y = 0\} = \frac{1}{2}$$

Here, $\mathbb{P}\{Y = 1\}$ denotes the probability of the event $Y = 1$ occurring. It should be noted that this random walk already satisfies properties $(i), (ii)$, and (iii) of standard Brownian motion if we constrain t to be a whole number given by the current timestep on our discrete walk.

Now, we will take some N and set the time increment to be $\Delta t = \frac{1}{N}$, defining a new random walk as

$$W_{k\Delta t} = a_N S_k$$

where a_N is a normalizing constant such that the variance of W_1 is still 1. Since the variance of S_1 is N, we see that $a_N = \frac{1}{\sqrt{N}}$ is our desired constant. An example with $N = 500$ can be seen in Figure 3. The details of the subsequent limiting process are omitted here; the intuition is simply that we reduce the time increments as N goes to infinity and thus create a continuous function satisfying the properties of standard Brownian motion, as desired.

Brownian motion also satisfies two important scaling properties:

Lemma 4.1. (Scaling Properties) ([2], Section 8.3) *Suppose X_t is a standard Brownian motion. Then*

(i) *If $a > 0$ and $Y_t = a^{-1/2} X_{at}$, then Y_t is a standard Brownian motion.*
(ii) *If X_t is a standard Brownian motion and $Y_t = tX_{1/t}$, then Y_t is a standard Brownian motion.*

Another important property of Brownian motion is the Markov property, which states that the movement of Brownian motion after any fixed time t is itself Brownian motion. In fact, Brownian motion satisfies a stronger version of this property, for which we need to define the notion of a stopping time. We say that a random variable T is a **stopping time** if for all t, observing the movement of the Brownian motion before t is enough to determine whether T has been reached. For example, if you are riding a bus down the street, the instruction "get off at the first red house you see" gives a stopping time but "get off at the last red house you see" does not (since there could always be another red house further along). As such,

$$T_x = \inf\{t \mid X_t = x\},$$

is a stopping time representing the first time that the Brownian motion attains a value of x. We say that Brownian motion satisfies the **strong Markov property**, since for all stopping times T, the process beyond time T

$$Y_t = X_{t+T} - X_T$$

is itself a Brownian motion independent of what X_t has done up until time T. [2]

We can use the strong Markov property to prove the following proposition.

Proposition 4.2. (Reflection Principle) *Suppose X_t is a Brownian motion with variance parameter σ^2 starting at a and $a < b$. Then for any $t > 0$,*

$$\mathbb{P}\{X_s \geq b \textit{ for some } 0 \leq s \leq t\} = 2\mathbb{P}\{X_t \geq b\},$$

where $\mathbb{P}\{X\}$ is the probability of event X occurring.

Proof. Let $T = T_b$ be a stopping time that denotes the first time that the Brownian motion reaches value b. Consequently, we note that

$$\mathbb{P}\{X_s \geq b \text{ for some } 0 \leq s \leq t\} = \mathbb{P}\{T \leq t\}$$

by continuity. Furthermore, since Brownian motion has a normal distribution, the probability that it takes on any particular value is 0. (Briefly, that this does not contradict the fact that Brownian motion has probability 1 of taking on some value in the reals has to do with the uncountability of the real numbers). Then since

$$\mathbb{P}\{T = t\} = \mathbb{P}\{X_t = b\} = 0,$$

we have that

$$\mathbb{P}\{T \leq t\} = \mathbb{P}\{T < t\}.$$

Now, we consider the event that $\{X_t \geq b\}$. Since X_t is normal with mean a and variance σ^2, we have

$$\mathbb{P}\{X_t \geq b\} = \int_b^\infty \frac{1}{\sqrt{2\pi t\sigma^2}} e^{-(x-a)^2/2\sigma^2 t} dx.$$

Also, since the event $\{X_t \geq b \mid T > t\}$ has probability 0 of occuring, we can write

$$\mathbb{P}\{X_t \geq b\} = \mathbb{P}\{T \leq t\}\mathbb{P}\{X_t \geq b \mid T \leq t\} = \mathbb{P}\{T < t\}\mathbb{P}\{X_t \geq b \mid T < t\}.$$

Now, suppose $T < t$. Then by the strong Markov property, if we let

$$Y_t = X_{t+T} - X_T = X_{t+T} - b,$$

then Y_t is itself a Brownian motion. Then $Y_{t-T} = X_t - b$ is a normal random variable with mean 0, and by the symmetry of the normal distribution we have that

$$\mathbb{P}\{X_t - b \geq 0 \mid T < t\} = 1/2.$$

Therefore, we can conclude that

$$\mathbb{P}\{T \leq t\} = 2\mathbb{P}\{X_t \geq b\} = 2\int_b^\infty \frac{1}{\sqrt{2\pi t\sigma^2}} e^{-(x-a)^2/2\sigma^2 t} dx.$$

□

5. Zero set of Brownian motion

One particularly interesting set that we can examine is the **zero set of Brownian motion,** defined as

$$Z = \{t \mid X_t = 0\}.$$

We follow Section 8.3 in [2] in exploring it.

A natural question to ask is whether standard Brownian motion will eventually return to its starting point after a certain period of time. As it turns out, the answer is yes, and this result can be proven using the zero set.

Theorem 5.1. *Standard Brownian motion will return to the origin an infinite number of times with probability* 1.

Proof. First, take $t > 1$. We want to find the probability that B_t returns to the origin between times 1 and t, or

$$\mathbb{P}\{Z \cap [1, t] \neq \emptyset\}.$$

Assume that $X_1 = b > 0$. The probability that $X_s = 0$ for some $1 \leq s \leq t$ given $X_1 = b$ is the same as the probability that $X_s \leq -b$ for some $0 \leq s \leq t-1$ without being given $X_1 = b$, which by symmetry is the probability that $X_s \geq b$ for $0 \leq s \leq t-1$. By the previously proven reflection principle, this implies that

$$\mathbb{P}\{X_s = 0 \text{ for some } 1 \leq s \leq t \mid X_1 = b\} = 2\int_b^\infty \frac{1}{2\pi(t-1)} e^{-x^2/2(t-1)} dx.$$

Again by symmetry, this probability is the same if $X_1 = -b$. We now seek to average the probability over all possible values of b. We can use the **probability density function** of Brownian motion $p_t(x, y)$ to do this, which yields the relative probability that X_t starting at x is arbitrarily close to y. Since $X_t - X_0$ is normal, has mean 0, and has variance t, we have

$$p_t(x, y) = \frac{1}{\sqrt{2\pi t}} e^{-(y-x)^2/2t},$$

where $-\infty < y < \infty$. Using this notation, we can write

$$\begin{aligned} \text{(5.2)} \quad & \mathbb{P}\{X_s = 0 \text{ for some } 1 \leq s \leq t\} \\ &= \int_{-\infty}^\infty p_1(0, b)\mathbb{P}\{X_s = 0 \text{ for some } 1 \leq s \leq t \mid X_1 = b\} db \\ &= 2\int_0^\infty \frac{1}{\sqrt{2\pi}} e^{-b^2/2} \left[2\int_b^\infty \frac{1}{\sqrt{2\pi(t-1)}} e^{-x^2/2(t-1)} dx\right] db. \end{aligned}$$

At this point, we can use the substitution $y = x/\sqrt{t-1}$ to simplify the integral to

$$4\int_0^\infty \int_{b(t-1)^{-1/2}}^\infty \frac{1}{2\pi} e^{-(b^2+y^2)/2} dy db,$$

which can be evaluated using polar coordinates. Specifically, we can use the outer integral to integrate over r and the inner integral to integrate over θ, yielding

$$\begin{aligned}(5.3)\quad 4\int_0^\infty \int_{\arctan((\sqrt{t-1})^{-1})}^{\pi/2} \frac{1}{2\pi} e^{-r^2/2} r\,d\theta\,dr \\ = 4\left(\frac{\pi}{2} - \arctan\frac{1}{\sqrt{t-1}}\right)\frac{1}{2\pi}\int_0^\infty re^{-r^2/2}dr \\ = 1 - \frac{2}{\pi}\arctan\frac{1}{\sqrt{t-1}}.\end{aligned}$$

Then this is the probability that Brownian motion will return to the origin at some point between 1 and t; in other words,

$$\mathbb{P}\{Z \cap [1,t] \neq \emptyset\} = 1 - \frac{2}{\pi}\arctan\frac{1}{\sqrt{t-1}}.$$

As t goes to infinity, the right-hand sum goes to 1; hence, Brownian motion is certain to eventually return to the origin. By the strong Markov property, since the movement after Brownian motion returns to the origin is itself a Brownian motion, this furthermore implies that it will return an infinite number of times. □

Another interesting property of Z is its fractal dimension. We will use a conception of dimension similar, but different from Hausdorff dimension, called box (or Minkowski) dimension.

Suppose A is a bounded set in $\mathbb{R}^d$. How many d-dimensional balls of diameter ϵ are required to cover A? For 1-dimensional lines, this number is on the order of ϵ^{-1} as ϵ becomes small. For 2-dimensional objects, like squares, this number is on the order of ϵ^{-2}. The **box dimension** of A, then, is the number D such that the number of ϵ-diameter balls required to cover A is on the order of ϵ^{-D} as ϵ becomes small.

As it turns out, box dimension and Hausdorff dimension satisfy

$$\dim_{\text{Hausdorff}}(E) \leq \dim_{\text{box}}(E)$$

for all sets E. Furthermore, for the sets we will consider in the rest of this paper, the two are equivalent.

Theorem 5.4. *The box dimension of Z is $1/2$ with probability 1.*

Proof. Consider $Z_1 = Z \cap [0,1]$ and $n \in \mathbb{N}$. We will try to cover Z_1 with balls of diameter $1/n$. For simplicity, we will divide the interval $[0,1]$ into n even subintervals as follows:

$$\left[\frac{k-1}{n}, \frac{k}{n}\right], \quad k = 1, 2, \cdots, n.$$

Note that this is similar to the alternative definition of box dimension discussed involving placing the set under consideration onto a grid. We now want to know how many of these intervals we need in order to cover the entire set. The probability of needing a specific interval is

$$P(k,n) = \mathbb{P}\left\{Z \cap \left[\frac{k-1}{n}, \frac{k}{n}\right] \neq \emptyset\right\}$$

By the scaling properties of Brownian motion, scaling our original process by $(k-1)/n$ (and adjusting the argument accordingly) still yields standard Brownian motion. Therefore, we have

$$P(k,n) = \mathbb{P}\left\{ Z \cap \left[1, \frac{k}{k-1}\right] \neq \emptyset \right\},$$

which we calculated above. Specifically,

$$P(k,n) = 1 - \frac{2}{\pi} \arctan \sqrt{k-1},$$

and the expected number of intervals to cover Z_1 is

$$\sum_{k=1}^{n} P(k,n) = \sum_{k=1}^{n} \left[1 - \frac{2}{\pi} \arctan \sqrt{k-1.}\right]$$

Recalling some trigonometry, we can use the Taylor series for $\arctan(1/t)$ to yield

$$\arctan x = \frac{\pi}{2} - \frac{1}{x} + O(\frac{1}{x^2}) \approx \frac{\pi}{2} - \frac{1}{x}$$

for large x, where $O(1/x^2)$ is a term on the order of $1/x^2$. Thus, we have

$$\sum_{k=1}^{n} P(k,n) \approx 1 + \sum_{k=2}^{n} \frac{2}{\pi\sqrt{k-1}} \approx \frac{2}{\pi} \int_1^n (x-1)^{-1/2} dx \approx \frac{4}{\pi}\sqrt{n}.$$

Therefore, it takes on the order of $\sqrt{n}$ intervals of length $1/n$ to cover Z_1, and the box dimension of Z_1 is $1/2$ with probability 1.

We then consider any $T > 0$ and define $\tilde{X}_s = T^{-1/2} X_{Ts}$. Then $\tilde{X}_s$ is standard Brownian motion, and

$$Z \cap [0,T] \text{ is similar to } \{s \in [0,1] \mid \tilde{B}_s = 0\}.$$

Since box dimension is invariant under similarities, and we have already proven the dimension of the right-hand set is $1/2$, we see that for every $T > 0$ we have with probability 1

$$\dim Z \cap [0,T] = 1/2$$

We now note that since Z can be represented as the countable union $\cup_{k=1}^{\infty}(Z \cap [0,k])$, we have

$$\dim(Z) = \sup_{k \in \mathbb{N}} \dim(Z \cap [0,k]).$$

Since we have already proven that each of the sets on the right has dimension $1/2$ with probability 1, and the intersection of countably many probability 1 events still has probability 1, we conclude that $\dim(Z) = 1/2$ with probability 1. □

The final interesting characteristic that we will investigate in Z is its topology. Using the scaling properties of Brownian motion in Lemma 4.1, we note that $Y_t = tX_{1/t}$ is also a standard Brownian motion. As time goes to infinity in the X process, time goes to 0 in the Y process. Then let Z_X and Z_Y denote the zero sets of X_t and Y_t respectively, noting that $Z = Z_X = Z_Y$. Since Z_X contains arbitrarily large values of t, Z_Y contains arbitrarily small values of t; therefore, if I is an interval containing 0, then $Z \cap I \neq \emptyset$.

The continuity of Brownian motion also implies that Z is closed: since $t_i \to t$ implies $X_{t_i} \to X_t$, we have that if a sequence of points $t_i \in Z$ converges to t, then $t \in Z$ as well. In the previous paragraph, we showed that 0 is not an isolated point in Z; i.e., there exist positive and decreasing points $t_i \in Z$ such that $t_i \to 0$. Though we will not prove this, it can be shown that none of the points of Z are isolated points and that Z contains no interval; any such set that is also non-empty and closed is topologically homeomorphic to the Cantor set. [2] [3]

6. Multi-dimensional Brownian motion

Suppose $X_t^1, X_t^2, \cdots, X_t^d$ are independent standard Brownian motions. Then we define the **standard d-dimensional brownian motion** as a vector-valued stochastic process

$$X_t = (X_t^1, X_t^2, \cdots, X_t^d).$$

The resulting stochastic process maintains continuity, randomness, mean 0, and the strong Markov property. Furthermore, we can calculate the probability density of X_t assuming $X_0 = x$ and denote it as

$$p_t(x, y) = \frac{1}{(2\pi t)^{d/2}} e^{-|y-x|^2/2t}.$$

We follow Sections 8.4, 8.5, and 8.6 of [2] in exploring multi-dimensional Brownian motion.

We now wish to ask whether multi-dimensional Brownian motion will always return to the origin, in a property we call **recurrence**. We have already shown that Brownian motion is recurrent in one-dimension; in other words, there are arbitrarily large times t such that $X_t = 0$.

First, let X_t be a standard d-dimensional Brownian motion. Let $0 < R_1 < R_2 < \infty$ and $B = B(R_1, R_2)$ be the annulus

$$B = \{x \in \mathbb{R}^d \mid R_1 < |x| < R_2.\}$$

We denote the boundary of the annulus by

$$\partial B = \{x \in \mathbb{R}^d \mid |x| = R_1 \text{ or } |x| = R_2\},$$

which can be thought of as the collection of points bordering B. We let $f(x) = f(x, R_1, R_2)$ be the probability that X_t starting at point x will reach the "outer" boundary $\{y \mid |y| = R_2\}$ before it reaches the "inner" boundary $\{y \mid |y| = R_1\}$,. Then if we let

$$\tau = \tau_{\partial B} = \inf\{t \mid X_t \in \partial B\},$$

we can write

$$f(x) = \mathbb{E}^x[g(X_t)]$$

where $g(y) = 1$ for $|y| = R_2$ and $g(y) = 0$ for $|y| = R_1$. As it turns out, by taking inspiration from the heat equation from physics and using some differential calculus that we will not get into here, we find that

$$f(x) = \phi(|x|) = \frac{\ln|x| - \ln R_1}{\ln R_2 - \ln R_1}, \quad d = 2,$$

$$f(x) = \phi(|x|) = \frac{R_1^{2-d} - |x|^{2-d}}{R_1^{2-d} - R_2^{2-d}}, \quad d \geq 3.$$

The symmetry of Brownian motion shows why $f(x)$ depends only on the distance of x from the origin. We can then ask what the probability is, in different dimensions, of Brownian motion returning arbitrarily close to the origin.

Take some $\epsilon > 0$ and let $R_1 = \epsilon$. As R_2 approaches infinity, what is the probability that Brownian motion started at some $x \in B$ reaches R_2 before R_1? In two-dimensions, this can be calculated by

$$\lim_{R_2 \to \infty} \mathbb{P}^x\{|X_t| = \epsilon \text{ before } |X_t| = R_2\} = \lim_{\epsilon \to 0} \left[1 - \frac{\ln|x| - \ln\epsilon}{\ln R_2 - \ln\epsilon}\right] = 0.$$

As a result, 2-dimensional Brownian motion will return arbitrarily close to the origin with probability 1. However, will it ever return exactly to the point 0, such that $X_t = 0$ for some t? If the answer is yes, there must be some R_2 such that X_t reaches 0 before reaching the circle of radius R_2 with positive probability. However, the probability of this event can be expressed as

$$\lim_{\epsilon \to 0} \mathbb{P}^x\{|X_t| = \epsilon \text{ before } |X_t| = R_2\} = \lim_{\epsilon \to 0} \left[1 - \frac{\ln|x| - \ln\epsilon}{\ln R_2 - \ln\epsilon}\right] = 0.$$

In other words, Brownian motion will never return to exactly 0, despite returning arbitrarily close to it an infinite number of times. We say that Brownian motion in two-dimensions therefore satisfies **neighborhood recurrence**, but not recurrence.

What about in $d \geq 3$? We can again take $\epsilon > 0$ and ask what the probability of Brownian motion never returning to the ball of radius ϵ is. Since $|x| > \epsilon$, we have

$$\lim_{R_2 \to \infty} \frac{\epsilon^{2-d} - |x|^{2-d}}{\epsilon^{2-d} - R_2^{2-d}} = 1 - \left(\frac{\epsilon}{|x|}\right)^{d-2} > 0.$$

Since this probability is greater than 0, the probability of returning to the ball is less than 1. As a result, the probability of returning an infinite number of times to the ball is 0, and eventually Brownian motion will escape any ball around the origin and go off to infinity. We say that Brownian motion in three dimensions or higher is therefore **transient**. Mathematician Shizuo Kakutani summarized these surprising results in a quote: *a drunk man will find his way home, but a drunk bird may get lost forever.*

Finally, we wish to consider what the fractal dimension of the path of multi-dimensional Brownian motion is. (The path of one-dimensional Brownian motion has dimension $3/2$. [4])

We will let X_t be a standard d-dimensional Brownian motion and A be

$$A = \{x \in \mathbb{R}^d \mid X_t = x \text{ for some } t\}.$$

In other words, A is the path of X_t. For ease, we consider the bounded set $A_1 = A \cap \{x \mid |x| \leq 1\}$; since the proof holds for all such bounded sets, we can take the limit as the bounds goes to infinity and prove it for A in general. We first look at the unit ball, $\{x \mid |x| \leq 1\}$, and create a minimal cover with balls of diameter ϵ. The number of balls required is on the order of ϵ^{-d}; how many of these balls are required to cover A_1?

In $d = 2$, the same argument that proves neighborhood recurrence can also prove that every open ball is visited by the Brownian motion. Therefore, every ball is needed, and the dimension of A is two.

In $d > 2$, we can use the calculations that proved transience to show that the probability that Brownian motion visits a ball of diameter $\epsilon/2$ centered around a point x (with $|x| > \epsilon/2$) is $(\epsilon/(2|x|))^{d-2}$. Therefore, when ϵ is small and $|x|$ is of order 1, the probability of visiting a generic ball of radius ϵ in the unit sphere is on the order of $c\epsilon^{d-2}$. Since each of the roughly ϵ^{-d} balls is chosen with probability on the order of $c\epsilon^{d-2}$, the total number of balls needed is on the order of $\epsilon^{d-2}\epsilon^{-d} = \epsilon^{-2}$, and the dimension of A is two.

Therefore, we conclude that the path of multi-dimensional Brownian motion (critically, regardless of the dimension) has dimension 2.

Acknowledgements

I would like to thank my mentor, Brin Harper, for her extensive support and knowledge without which this paper would not have been possible. I would also like to thank the many volunteers and lecturers who created a mathematically enriching experience this summer, in particular Jinwoo Sung for inspiring parts of this paper. Finally, I would like to thank Peter May for organizing and running this REU.

7. Bibliography

References

[1] Rami Shakarchi, Elias M. Stein. Real Analysis: Measure Theory, Integration, and Hilbert Spaces. Princeton University Press. 2009.
[2] Gregory F. Lawler. Introduction to Stochastic Processes. Chapman and Hall/CRC. 2006.
[3] https://galton.uchicago.edu/~lalley/Courses/313/BrownianMotionCurrent.pdf
[4] Peter Mörters, Yuval Peres. Brownian Motion. Cambridge University Press. 2010.
[5] https://en.wikipedia.org/wiki/Sierpinski_carpet
[6] https://en.wikipedia.org/wiki/Menger_sponge
[7] https://en.wikipedia.org/wiki/Cantor_set#Cantor_dust
[8] Images sourced without modification from Wikimedia Commons, originally published by Niabot and WalkingRadiance, under the Creative Commons Attribution 3.0 and Attribution-ShareAlike 4.0 International licenses.
[9] Code used to run simulations can be found at this Github link.

Appendix A. Proofs of lemmas

Lemma 3.2. There exists a closed ball B such that $S_k(B) \subset B$ for all $k = 1, \cdots, m$.

Proof. Note that if S is a similarity with ratio r, we can use the definition and the triangle inequality to yield

$$\begin{aligned} |S(x)| &\leq |S(x) - S(0)| + |S(0)| \\ &\leq r|x| + |S(0)|. \end{aligned}$$

If R is the radius of our desired ball B (which we will center on the origin), we want an R such that if $|x| \leq R$ (equivalently, if $x \in B$), then $|S(x) \leq R|$ ($S(x) \in B$). Then by the above inequality, we want to satisfy $rR + |S(0)| \leq R$, or $R \geq |S(0)|/(1-r)$. In this way, we can find a B_k for each S_k such that $S_k(B_k) \subset B_k$, and let B be the B_k with the largest radius R. Thus, $S_k(B) \subset B$ for all k. $\square$

Lemma 3.3. The distance function dist on compact sets in $\mathbb{R}^d$ satisfies the following properties:

(i) $\mathrm{dist}(A, B) = 0$ if and only if $A = B$.
(ii) $\mathrm{dist}(A, B) = \mathrm{dist}(B, A)$.
(iii) $\mathrm{dist}(A, B) \leq \mathrm{dist}(A, C) + \mathrm{dist}(C, B)$.
(iv) $\mathrm{dist}(\tilde{S}(A), \tilde{S}(B)) \leq r\mathrm{dist}(A, B)$.

Proof. First, note that if $A \neq B$, then without loss in generality there must exist some point x such that $x \in A \backslash B$. Since B is compact, we can let $\delta = d(x, B)/2 > 0$ and find that $B \not\subset A^\delta$, so $\operatorname{dist}(A, B) > \delta > 0$. To prove the other direction, we can find that if $\operatorname{dist}(A, B) \neq 0$, then without loss in generality there exists some δ and $x \in A$ such that $x \notin B^\delta \supset B$, which implies that $A \neq B$. Thus, (i) is proven.

To prove (ii), we use the symmetry of the definition.

To prove (iii), we see that if $A \subset C^{\delta_1}$ and $C \subset B^{\delta_2}$, then for any $a \in A$, there is a $b \in B$ and $c \in C$ such that $d(a, c) < \delta_1$ and $d(c, b) < \delta_2$. Thus, by the triangle inequality, $d(a, b) < \delta_1 + \delta_2$, so $A \subset B^{\delta_1+\delta_2}$. Since the inverse can be similarly proven, we see that (iii) is satisfied.

To loosely prove (iv), we use the fact that similarities are compositions of rotations, translations, and dilations. Hausdorff distance can be shown to be rotation and translation-invariant, while dilations by r can only dilate the distance by a factor of r. □

Lemma 3.4. There exists a unique non-empty compact set F such that

$$F = S_1(F) \cup \cdots \cup S_m(F).$$

Proof. First, we take B as in Lemma 3.2, and let $F_k = \tilde{S}^k(B)$ as defined above. Since B is compact and non-empty, each F_k is as well, and sinec $\tilde{S}(B) \subset B$, we have $F_k \subset F_{k-1}$. Then we let

$$F = \bigcap_{k=1}^{\infty} F_k,$$

noting that F is compact, non-empty, and that $\tilde{S}(F) = \cap_{k=2}^{\infty} F_k = F$. We can also prove that F is unique. Suppose G is another compact non-empty set such that $\tilde{S}(G) = G$. Then by Lemma 3.3, $\operatorname{dist}(F, G) \leq r\operatorname{dist}(F, G)$, and since $r < 1$, this forces $\operatorname{dist}(F, G) = 0$, so $F = G$. Thus, the proof is complete. □

Lemma 3.8. Suppose B is a ball in the covering $\mathcal{B}$ that satisfies

$$r^l \leq \operatorname{diam} B \leq r^{l-1} \quad \text{for some } l \leq k.$$

Then B contains at most cm^{k-l} vertices of the k^{th} generation.

Proof. Let v be a vertex of the k^{th} generation such that $v \in B$, and let $\mathcal{O}(v)$ be its corresponding open set. Then properties (i) and (ii) above guarantee that there exists some fixed dilate $B*$ of B such that $\mathcal{O}(v) \subset B*$ and $B*$ contains the open set of generation l that contains $\mathcal{O}(v)$. We note that each open set in generation l has volume approximately equal to r^{dl} by property (ii) above. Since $B*$ has volume cr^{dl} (since r^{dl} is approximately the volume of B), $B*$ can contain at most c open sets of generation l, and therefore at most cm^{k-l} open sets of generation k. Thus, B can also contain at most cm^{k-l} vertices of generation k, and the lemma is proved. □