HAUSDORFF DIMENSION OF RANDOM FRACTALS

SHRAVAN HARIBALARAMAN

Abstract. Self-similar fractals are of particular interest to geometric measure theorists because their properties (such as Hausdorff dimension) are easy to analyze, due to Hutchinson's [1] classical construction. This paper explores how stochasticity and randomness may be introduced into the fractal generating process, and investigates ways to extend and simplify previous constructions by Graf [2] and Maudlin-Williams [3] to determine the Hausdorff dimension of random fractals generated by a stochastic process. Our primary result highlights an elegant parallel between the deterministic and random cases.

Contents

1. Introduction

The study of fractals often begins with the most rudimentary class of fractals: self-similar ones. Roughly speaking, a fractal is self-similar if it is the union of scaled down copies of itself. The self-similarity of a set makes it easier for us to calculate its dimension and measure; however, it is unfortunately the case that self-similarity is an incredibly strong condition to impose on a fractal, and most sets with complex, irregular fractal structures are not self-similar.

One way to extend the tools we have to understand self-similar sets is by introducing stochasticity into the fractal generating process. Section 2 of this paper formalizes self-similarity and stochastic self-similarity through the theory of function systems, while Section 3 introduces Hausdorff dimension and introduces tools to analyze the dimension of self-similar fractals. Sections 4 and 5 extend these tools to the random/stochastic case to determine the dimension of random fractals, using several proofs and techniques introduced by Graf [2] and Maudlin-Williams

Date: August 22, 2025.

[3]. Section 4 contains Theorem 4.3, the paper's primary result (a variant of which was originally shown by Maudlin-Williams [3]): that the Hausdorff dimension of a typical random fractal is the unique value that solves a simple expectation equation.

2. (Stochastic) Self-Similarity

In this section we define self-similarity and explain how, using contraction maps, we may generate fractals that are self-similar and 'stochastically' self-similar.

2.1. **Contractions, GFSs, and the Hausdorff Metric.** Roughly speaking, a self-similar set is a set that is comprised of smaller copies of itself, and therefore has complex yet well-behaved structures at arbitrarily small scales. We introduce fractals through a theory of generalized function systems. Throughout this section, (X, d) is a compact metric space, unless otherwise specified.

Definition 2.1. For a function $f : X \to X$, the *Lipschitz constant* of f is defined as

$$\mathrm{Lip}(f) = \sup_{x,y \in X} \frac{d(f(x), f(y))}{d(x,y)}.$$

If $\mathrm{Lip}(f) < \infty$, f is called *Lipschitz continuous.* If $\mathrm{Lip}(f) < 1$, f is called a *contraction.* If there exists $r \geq 0$ such that for all $x, y \in X$, $d(f(x), f(y)) = rd(x, y)$, f is called a *similitude.* If $0 \leq r < 1$, f is called a *contractive similitude.*

We now build up the theory of function systems. Until Section 4, note that N is some fixed natural number. Section 5 deals with some possible realizations of N.

Definition 2.2. For $N \in \mathbb{N}$, let $[N] = \{1, 2, \ldots, N\}$. Let $[N]^*$ denote the set of all non-empty finite sequences, and $[N]^{\mathbb{N}}$ be the set of all infinite sequences, in $[N]$.

- For $\alpha = (n_1, \ldots, n_k) \in [N]^*$, $|\alpha| = k$ is the *length* of σ.
- For $\sigma \in [N]^* \cup [N]^{\mathbb{N}}$, define $\sigma \mid j = (n_1, \ldots, n_j)$ to be the *j-th curtailment* of σ.
- For $\sigma_1 = (n_1, \ldots, n_k) \in [N]^*$ and $\sigma_2 = (m_1, \ldots, m_l) \in [N]^*$, we define the *concatenation* of σ_1 and σ_2 to be $\sigma_1 \bullet \sigma_2 = (n_1, \ldots, n_k, m_1, \ldots, m_l)$.

It is helpful to visualize the above as a countable branching tree wherein each node has N offsprings. The next step is to assign to each node a contraction.

Definition 2.3. Let $\mathscr{S}(X)$ denote the set of all contractive similitudes on X. Then, a *generalized function system* (GFS) is a map $F : [N]^* \to \mathscr{S}(X)$. We will denote the contraction $F(\sigma)$ by F_σ for convenience, and shorten $F_{(n)}$ to F_n, where (n) is the sequence of length 1.

Definition 2.4. The *GFS attractor* of a GFS F is the set

$$K_F = \bigcap_{q \in \mathbb{N}} \bigcup_{|\sigma| = q} F_{\sigma|1} \circ \cdots \circ F_{\sigma|q}(X).$$

Proposition 2.5. *If F is a GFS, its GFS attractor K_F is compact. Furthermore, if, for all $q \geq 1$, there exists σ such that $|\sigma| = q$ and $Lip(F_{\sigma|1} \circ \cdots \circ F_{\sigma|q}) > 0$, then K_F is non-empty.*

Proof. For $q \geq 1$ let $K_q = \bigcup_{|\sigma|=q} F_{\sigma|1} \circ \cdots \circ F_{\sigma|q}(X)$. For each σ and each $i \geq q$, $F_{\sigma|i}$ is a similitude, so their composition is a similitude, implying that the set $F_{\sigma|1} \circ \cdots \circ F_{\sigma|q}(X)$ is similar to X, and therefore also compact. There are at most N^q distinct sequences σ with length q, so K_q is a finite union of compact sets.

Now, for any σ with $|\sigma| = q+1$,

$$F_{\sigma|1} \circ \cdots \circ F_{\sigma|(q+1)}(X) = F_{\sigma|1} \circ \cdots \circ F_{\sigma|q}(F_{\sigma|(q+1)}(X)) \subset F_{\sigma|1} \circ \cdots \circ F_{\sigma|q}(K),$$

since $F_{\sigma|(q+1)}(X) \subset X$. This implies that $F_{\sigma|1} \circ \cdots \circ F_{\sigma|(q+1)}(X) \subset K_q$, and therefore, $K_{q+1} \subset K_q$. $K_F = \bigcap_q K_q$ is therefore a countable intersection of nested compact sets, which must be compact. The additional assumption that there is at least one non-negative Lipschitz constant at each $q \geq 1$ implies that K_q is non-empty for each q; thus, their intersection must be non-empty. □

When dealing with deterministic self-similarity, we are almost always interested in a particular case of GFS, known as an 'iterated function system'.

Definition 2.6. An *iterated function system* (IFS) is a GFS where, for all $n \leq N$ and $\sigma \in [N]^*$, $F_n = F_{\sigma \bullet n}$. The GFS attractor of an IFS is called its *IFS attractor*. A set is called *self-similar* if it is the IFS attractor of some IFS.

The following equivalence is immediate from the definition of an IFS.

Proposition 2.7. *Let F be an IFS. Then K_F is its IFS attractor iff K_F is compact and*

$$K_F = F_1(K_F) \cup \cdots \cup F_N(K_F).$$

Example 2.8. Say $N = 2$, $X = [0,1]$, $F_1(x) = \frac{1}{3}(x)$, and $F_2(x) = \frac{1}{3}(x) + \frac{2}{3}$. The IFS attractor K_F of this IFS is none other than the middle-thirds Cantor set, and the sequence $\{K_q\}$ from Proposition 2.5 is merely the decreasing sequence of sets used in the classical construction of the Cantor set by 'removing' the middle thirds.

Proposition 2.9. *Let F be an IFS with attractor K_F. For $\sigma = (n_1, \ldots, n_q) \in [N]^q$, denote $K_F^\sigma = F_{\sigma|1} \circ \cdots \circ F_{\sigma|q}(K_F) = F_{n_1} \circ \cdots \circ F_{n_q}(K_F)$. Then, if the Lipschitz constants $\mathrm{Lip}(F_1), \ldots, \mathrm{Lip}(F_N)$ are strictly positive, for every $\tau \in [N]^{\mathbb{N}}$ the set*

$$\bigcap_{q=1}^{\infty} K_F^{\tau|q}$$

is a singleton.

Proof. The set in question is a countable intersection of nested compact sets; furthermore, $\mathrm{diam}(K_F^{\tau|q}) \leq \max_{1 \leq i \leq N} L_i^q \cdot \mathrm{diam}(K_F) \to 0$ as $q \to \infty$; so the intersection is a point. We will denote this point x_τ for convenience. □

The next natural question is: does the set K_F from Proposition 2.7 always exist? The answer is positive, as shown in Theorem 2.12.

Definition 2.10. Let $\mathscr{K}(X)$ denote the set of non-empty compact subsets of X and, for $x \in X$ and $K \subset X$, let $\mathrm{dist}(x, K) = \inf_{y \in K} d(x, y)$.

The map $\rho : \mathscr{K}(X) \times \mathscr{K}(X) \to [0, \infty)$, where

$$\rho(K_1, K_2) = \max\{\sup_{x \in K_1} \mathrm{dist}(x, K_2), \sup_{y \in K_2} \mathrm{dist}(y, K_1)\},$$

is called the *Hausdorff metric* on $\mathscr{K}(X)$.

ρ truly is a metric on $\mathscr{K}(X)$ - this is easy to check. What is harder to check, but is a crucial fact, is the following statement, proved in [1]. In fact, it holds more generally whenever X is complete:

Lemma 2.11. *$(\mathscr{K}(X), \rho)$ is a complete metric space.*

We come to the most important theorem thus far, that provides an alternative characterization of the IFS attractor of an IFS F.

Theorem 2.12. *Let F be an IFS and define the map $G_F : \mathscr{K}(X) \to \mathscr{K}(X)$ by*

$$G_F(K) = F_1(K) \cup \cdots \cup F_N(K).$$

Then, if $K \subset X$ is compact, $G_F^q(K) \to K_F$ in the Hausdorff metric, where K_F is the IFS attractor of F.

Proof. Define the constant $L = \max_i \mathrm{Lip}(F_i)$, and consider two arbitrary sets $K_1, K_2 \in \mathscr{K}(X)$.

Fixing $x \in K_1$, we see that for all $1 \leq i \leq N$,

$$\mathrm{dist}(F_i(x), G_F(K_2)) \leq \mathrm{dist}(F_i(x), F_i(K_2)) \leq L_i \cdot \mathrm{dist}(x, K_2) \leq L \cdot \mathrm{dist}(x, K_2),$$

and thus,

$$\sup_{y \in G_F(K_1)} \mathrm{dist}(y, G_F(K_2)) \leq L \cdot \sup_{x \in K_1} \mathrm{dist}(x, K_2).$$

Clearly the same holds when K_1 and K_2 are swapped; and so,

$$\rho(G_F(K_1), G_F(K_2)) \leq L \cdot \rho(K_1, K_2) \implies \mathrm{Lip}(G_F) \leq L < 1.$$

Therefore, G_F is a contraction on $\mathscr{K}(X)$, a complete metric space. Per Banach's contraction mapping theorem, it has a unique fixed point $K_F \in \mathscr{K}(X)$. Furthermore, for any compact $K \subset X$, $G_F^q(K) \to K_F$ as $q \to \infty$, concluding the proof. $\square$

Remark 2.13. Note that this proof did not use the fact that $F_1, \ldots, F_N$ were similitudes. The same result actually holds true if they are just general contraction mappings.

An IFS F is completely characterized by the N-tuple $(F_1, \ldots, F_N) \in \mathscr{S}(X)^N$. So, Theorem 2.12 provides, in some sense, a map between N-tuples of contractive similitudes and the self-similar compact sets they generate.

2.2. **Generating Random Fractals.** So far stochasticity hasn't come into the picture; but the general idea for generating random fractals is related, with the natural inclusion of some measure theory.

This process can be laid out in three steps, beginning with a probability measure on N-tuples of contractions and then ending with a probability measure on the set of GFS attractors.

2.2.1. *Step 1: A Probability On $\tilde{\mathscr{S}}(X)^N$.* Begin with an arbitrary probability measure μ on the set $\tilde{\mathscr{S}}(X)^N \subset \mathscr{S}(X)^N$ of N-tuples of contractive similitudes that follow a crucial 'disjointness' property: for $(F_1, \ldots, F_N) \in \mathscr{S}(X)^N$, whenever $i \neq j$ we have $\mathrm{int}(F_i(X)) \cap \mathrm{int}(F_j(X)) = \emptyset$. The following is immediate:

Proposition 2.14. $2^{\tilde{\mathscr{S}}(X)^N}$*, the power set of $\tilde{\mathscr{S}}(X)^N$, is a σ-algebra, so μ is well-defined.*

Note that because μ is a probability, we can take arbitrary products of μ with itself. Consider $(\tilde{\mathscr{S}}(X)^N)^{[N]^*}$, the set of all maps from $[N]^*$ to $\tilde{\mathscr{S}}(X)^N$. In Step 2, we will interpret an element of this set as a 'pulled-back GFS', wherein each $\sigma \in [N]^*$ is mapped to the N-tuple $(F_{\sigma \bullet 1}, \ldots, F_{\sigma \bullet N})$. This set is endowed with the product measure $\mu^{[N]^*}$.

2.2.2. *Step 2: A Probability On $\Delta(X)$.* We use μ to develop a probability measure on a well-chosen set of GFSs.

Definition 2.15. Let $\Gamma(X)$ be the set of all GFSs on X, and $\Delta(X) \subset \Gamma(X)$ be the set of all GFSs on X satisfying the following condition:

For all $\sigma \in [N]^*$ and $n_1, n_2 \in [N]$ where $n_1 \neq n_2$,

$$\operatorname{int}(F_{\sigma|1} \circ \cdots \circ F_\sigma \circ F_{\sigma \bullet n_1}(X)) \cap \operatorname{int}(F_{\sigma|1} \circ \cdots \circ F_\sigma \circ F_{\sigma \bullet n_2}(X)) = \emptyset.$$

$\Delta(X)$ is referred to as the set of *disjoint GFSs* on X.

We suppress the argument and just use Γ and Δ from here onward. The following is, again, immediate:

Proposition 2.16. $\Sigma = 2^\Delta$*, the power set of* Δ*, is a* σ*-algebra.*

Here, we run into no problems with non-measurability because Σ is simply the product of the σ-algebras $2^{\tilde{\mathscr{S}}(X)^N}$ previously defined; so every disjoint GFS is measurable. So, what we now do is identify each disjoint GFS with an N-tuple of contractions and an element of $(\tilde{\mathscr{S}}(X)^N)^{[N]^*}$ in a natural way.

Definition 2.17. The *pushforward map* $\varphi : \tilde{\mathscr{S}}(X)^N \times (\tilde{\mathscr{S}}(X)^N)^{[N]^*} \to \Delta$ is the map that sends $((g_1, \dots, g_N), f) \in \tilde{\mathscr{S}}(X)^N \times (\tilde{\mathscr{S}}(X)^N)^{[N]^*}$ to the GFS F, where $F_{(n)} = g_n$ and for all $\sigma \in [N]^*$, $f(\sigma) = (F_{\sigma \bullet 1}, \dots, F_{\sigma \bullet N})$.

Proposition 2.18. φ *is measurable and bijective.*

The above is easily checked. Now, the domain of φ is endowed with the measure $\mu \times \mu^{[N]^*}$; thus, we may define a measure ν on Δ as the pushforward measure of $\mu \times \mu^{[N]^*}$ under the map φ. That is, for $A \subset \Delta$, $\nu(A) = (\mu \times \mu^{[N]^*})(\varphi^{-1}(A))$. This is denoted by $\nu = \varphi_{\#\mu \times \mu^{[N]^*}}$.

Note that since φ is bijective, we may take the dual approach and equivalently define ν by the fact that $(\mu \times \mu^{[N]^*})(B) = \nu(\varphi(B))$ for $B \subset \tilde{\mathscr{S}}(X)^N \times (\tilde{\mathscr{S}}(X)^N)^{[N]^*}$.

2.2.3. *Step 3: A Probability On $\mathscr{K}(X)$.* Now, we can use the probability ν on Δ to define a final measure on the set of attractors of disjoint GFSs. This is, again, done in the natural way. Define the map $u : \Delta \to \mathscr{K}(X)$ by $u(F) = K_F$, where K_F is the attractor of F. Then, we can define $\mathbb{P} = u_{\#\nu}$ as the pushforward measure of ν under the map u.[1]

That is, for $A \subset \mathscr{K}(X)$, $\mathbb{P}(A) = \nu(u^{-1}(A))$. There are no guarantees that u is bijective, but what is true is that if $\tilde{\mathscr{S}}(X)$ is endowed with the topology of pointwise convergence, u is continuous and therefore measurable, so $\mathbb{P}$ is well-defined. Clearly, the set of all disjoint GFS attractors has full $\mathbb{P}$-measure.

Definition 2.19. A *stochastic IFS* is a double (F, μ), where μ is a probability on $\tilde{\mathscr{S}}(X)^N$ and F is a random map from $[N]^*$ to $\tilde{\mathscr{S}}(X)$, with the probability distribution governed by the measure $\nu = \varphi_{\#\mu \times \mu^{[N]^*}}$. A set $K \in \mathscr{K}(X)$ is called μ*-self-similar* if it lies in the support of the measure $\mathbb{P}$ constructed in Step 3 above; K_F refers to the random set with probability distribution $\mathbb{P}$.

Remark 2.20. This procedure tells us how we can generate well-behaved stochastically self-similar sets. Given a probability distribution μ on $\tilde{\mathscr{S}}(X)^N$, choose one N-tuple and then $[N]^*$ N-tuples of contractive similitudes. This is identified with a unique GFS and thus its attractor.

[1] $\mathbb{P}$ should be distinguished from P, which will be used more generally to represent the probability of any arbitrary event.

3. Hausdorff Measure and Dimension

In this section we introduce the properties of interest in this paper: Hausdorff measure and dimension, and we show how they can be calculated for self-similar sets in $\mathbb{R}^d$.

3.1. **Constructing Hausdorff Measure and Dimension.**

There are many equivalent definitions and constructions of the Hausdorff measure; we provide one adapted from Matilla [4].

Definition 3.1. For $\varepsilon > 0$ and $s \in [0, \infty)$, the *s-dimensional Hausdorff ε-measure* of a set $A \subset \mathbb{R}^d$ is given by

$$\mathscr{H}^s_\varepsilon(A) = \inf \left\{ \sum_{i=1}^{\infty} \operatorname{diam}(U_i)^s : A \subset \bigcup_{i=1}^{\infty} U_i, \ \operatorname{diam}(U_i) < \varepsilon \right\},$$

where $\operatorname{diam}(U) = \sup\{\|x - y\| : x, y \in U\}$ is the diameter of U.

While $\mathscr{H}^s_\varepsilon$ is monotonic and $\mathscr{H}^s_\varepsilon(\emptyset) = 0$, it is not countably additive... yet.

Definition 3.2. The *s-dimensional Hausdorff measure* of a set $A \subset \mathbb{R}^d$ is defined by

$$\mathscr{H}^s(A) = \lim_{\varepsilon \to 0^+} \mathscr{H}^s_\varepsilon(A).$$

This limit always exists since $\mathscr{H}^s_\varepsilon(A)$ is easily seen to be a non-increasing function of ε. From this definition, the following facts are easily checked; proofs are provided by Matilla [4].

Proposition 3.3. *$\mathscr{H}^s$ is a well-defined measure. Additionally,*

1. *All Borel sets are $\mathscr{H}^s$-measurable.*
2. *$\mathscr{H}^s$ is Borel regular. That is, for all $X \subset \mathbb{R}^d$, there exists a $B \supseteq A$ such that B is Borel and $\mathscr{H}^s(B) = \mathscr{H}^s(A)$.*
3. *$\mathscr{H}^s$ is translation invariant.*
4. *$\mathscr{H}^s(kA) = k^s \mathscr{H}^s(A)$, where $kX = \{k \cdot x : x \in X\}$.*

Given $A \subset \mathbb{R}^d$, the following theorem tells us that for all small enough s, $\mathscr{H}^s(A) = \infty$ and for all large enough $\mathscr{H}^s(A) = 0$.

Theorem 3.4. *Let $A \subset \mathbb{R}^d$ and $s, t \in [0, \infty)$ so that $s < t$. Then,*

$$\mathscr{H}^t(A) > 0 \implies \mathscr{H}^s(A) = \infty.$$

Proof. For fixed $\varepsilon > 0$, consider any countable open cover $\{U_i\}$ of A such that $\operatorname{diam}(U_i) < \varepsilon$; consequently, $\operatorname{diam}(U_i)^{s-t} > \varepsilon^{s-t}$. Then,

$$\sum_{i=1}^{\infty} \operatorname{diam}(U_i)^s = \sum_{i=1}^{\infty} \operatorname{diam}(U_i)^t \cdot \operatorname{diam}(U_i)^{s-t} > \varepsilon^{s-t} \cdot \sum_{i=1}^{\infty} \operatorname{diam}(U_i)^t \geq \varepsilon^{s-t} \mathscr{H}^t_\varepsilon(A).$$

Thus, $\mathscr{H}^s_\varepsilon(A) > \varepsilon^{s-t} \mathscr{H}^t_\varepsilon(A)$. Letting $\varepsilon \to 0^+$, we obtain the desired result. $\square$

This result, along with its contrapositive, implies that for each set X there is one 'interesting' value of s for which the s-dimensional Hausdorff measure of X does not collapse or blow up.[2] This yields the following definition:

[2] Of course, for this value of s, the measure might still be zero or infinite.

Definition 3.5. The *Hausdorff dimension* of a set $A \subset \mathbb{R}^d$ is defined as

$$\dim_{\mathscr{H}}(A) = \sup\{s : \mathscr{H}^s(A) = \infty\}.$$

Remark 3.6. We could alternatively define the Hausdorff dimension as

$$\dim_{\mathscr{H}}(A) = \inf\{s : \mathscr{H}^s(A) = 0\}.$$

Clearly these definitions are equivalent.

Often when dealing with fractals we care about their Hausdorff dimension rather than their Hausdorff measure because the measure can be renormalized easily once the dimension is known.

3.2. **The Hausdorff Dimension of Self-Similar Sets.** Finding the Hausdorff dimension of a compact self-similar set in $\mathbb{R}^d$ is (relatively) easy, especially since we have built up all the machinery to do so in Section 2. We need just one more definition:

Definition 3.7. An IFS F on a compact set $K \subset \mathbb{R}^d$ is said to satisfy the *open set condition* (OSC) if there exists a nonempty bounded open set $U \subset \mathbb{R}^d$ such that $U \subset G_F(U)$ (where G_F is the function defined in Theorem 2.12) and the sets $F_i(U)$ are all disjoint.

Example 3.8. The IFS that defines the Cantor set satisfies the open set condition; letting $U = (0, 1)$ makes this clear.

Theorem 3.9. *Let K_F be the IFS attractor of F, where F satisfies the OSC; and denote $L_i = \mathrm{Lip}(F_i)$ for $i = 1, \ldots, N$. Then, the Hausdorff dimension of K_F is s, where s is the unique real number such that*

$$\sum_{i=1}^{N} L_i^s = 1.$$

Proof. We break the proof into two parts.

Part I. Suppose $\sum_{i=1}^{N} L_i^s = 1$ and fix $k \in \mathbb{N}$. Note that

$$K_F = \bigcup_{i=1}^{N} F_i(K_F) = \bigcup_{i=1}^{N} \bigcup_{j=1}^{N} F_i \circ F_j(K_F),$$

and as we iterate this, we obtain

$$K_F = \bigcup_{|\sigma|=q} K_F^\sigma$$

for any given $q \in \mathbb{N}$. In particular, $K_F \subset \bigcup_{|\sigma|=k} K_F^\sigma$. ($K_F^\sigma$ was introduced in Proposition 2.9).

Now, fix $\varepsilon > 0$. Letting $L = \max_i L_i$, we choose q_0 so large that $L^{q_0} < \frac{\varepsilon}{\mathrm{diam}(K_F)}$. Then, for $\sigma \in [N]^{q_0}$,

$$\mathrm{diam}(K_F^\sigma) = \mathrm{diam}(K_F) \cdot \prod_{i=1}^{q_0} L_i \leq L^{q_0} \cdot \mathrm{diam}(K_F) < \varepsilon.$$

Thus,

$$\mathscr{H}_\varepsilon^s(K_F) \leq \sum_{|\sigma|=q_0} \mathrm{diam}(K_F^\sigma)^s = \sum_{i=1}^{d} \sum_{|\sigma|=q_0-1} L_i^s \cdot \mathrm{diam}(K_F^\sigma)^s = \sum_{|\sigma|=q_0-1} \mathrm{diam}(K_F^\sigma)^s,$$

since $\sum_{i=1}^N L_i^s = 1$. Iterating this down to $|\sigma| = 1$, we simply get

$$\mathscr{H}_\varepsilon^s(K_F) \leq \sum_{i=1}^N L_i^s \cdot \operatorname{diam}(K_F)^s = \operatorname{diam}(K_F)^s.$$

Since ε was chosen arbitrarily, the same holds in the limit, and thus we obtain $\mathscr{H}^s(K_F) \leq \operatorname{diam}(K_F)^s < \infty$.

Part II. For $\sigma \in [N]^q, \tau \in [N]^{\mathbb{N}}$, we will say σ *precedes* τ (denoted $\sigma \prec \tau$) if $\tau \mid q = \sigma$. For $\sigma \in [N]^*$, denote $[N]_\sigma = \{\tau \in [N]^{\mathbb{N}} : \sigma \prec \tau\}$.

We introduce a measure λ on $[N]^{\mathbb{N}}$ as follows: for $\sigma = (n_1, \ldots, n_q)$, we will define $\lambda([N]_\sigma) = (L_{n_1} L_{n_2} \ldots L_{n_q})^s$. It is easy to see that $\lambda([N]^{\mathbb{N}}) = 1$. Furthermore, $\{[N]_\sigma\}_{\sigma \in [N]^*}$ is a collection of open sets that generates the Borel σ-algebra on $[N]^{\mathbb{N}}$ (under the product of discrete topologies), and so by Carathéodory's Extension Theorem, λ can be extended to a probability on all of $[N]^{\mathbb{N}}$.

We use λ to generate a measure λ^* on subsets of $\mathbb{R}^d$: for $A \subset \mathbb{R}^d$, define

$$\lambda^*(A) = \lambda(\{\tau \in [N]^{\mathbb{N}} : x_\tau \in A \cap K_F\}),$$

where x_τ is defined in Proposition 2.9. This implies that $\lambda^*(K_F) = 1$, and for any $\sigma \in [N]^q$,

$$\lambda^*(K_F^\sigma) = \lambda(\{\tau \in [N]^{\mathbb{N}} : x_\tau \in K_F^\sigma\}) = \lambda(\{[N]_\sigma\}) = (L_{n_1} L_{n_2} \ldots L_{n_q})^s.$$

Let U be a non-empty bounded open set such that $G_F(U) \subset U$ and the sets $F_i(U)$ are disjoint. By the OSC, U exists; furthermore, $G_F^q(\overline{U}) \to K_F$ as $q \to \infty$.

Now choose $r \in (0,1)$ and define the following subset $\omega_r \subset [N]^*$: for each $\tau = \{n_q\}_{q=1}^\infty \in [N]^{\mathbb{N}}$, curtail τ at n_q, where q is the first index for which

$$\min_{1 \leq i \leq N} L_i \cdot r \leq (L_{n_1} \ldots L_{n_q}) \leq r,$$

and add $\tau \mid q$ to ω_r. Similarly to Proposition 2.9, for $\sigma = (n_1, \ldots, n_q)$ define $U_F^\sigma = F_{n_1} \circ \cdots \circ F_{n_q}(U)$. Since $U_F^1, \ldots, U_F^N$ are disjoint and nested within U, it holds that $U_F^{(\tau|q)\bullet 1}, \ldots, U_F^{(\tau|q)\bullet N}$ are disjoint and nested within U; from this it is observed that $\{U_F^\sigma : \sigma \in \omega_r\}$ is a disjoint collection of subsets of U. Since, for each $\tau \in [N]^{\mathbb{N}}$, there exists $\sigma \in \omega_r$ such that $\sigma \prec \tau$, $x_\tau \in K_F^\sigma$, and so

$$K_F = \bigcup_{\tau \in [N]^{\mathbb{N}}} x_\tau \subset \bigcup_{\sigma \in \omega_r} K_F^\sigma \subset \bigcup_{\sigma \in \omega_r} \overline{U_F^\sigma}.$$

(The last inclusion is a consequence of the fact that $G_F^q(\overline{U}) \to K_F$.)

Since U is open it contains a ball of small radius a; and since it's bounded it is contained within a ball of larger radius b. Then, for $\sigma = (n_1, \ldots, n_q) \in \omega_r$, it follows that U_F^σ is contains a ball of radius $((\min_{1 \leq i \leq N} L_i) \cdot a \cdot r) \leq (a \cdot L_{n_1} \ldots L_{n_q})$, and it's contained in a ball of radius $(b \cdot r) \geq (b \cdot L_{n_1} \ldots L_{n_q})$.

For a ball B with radius r, let $\omega_{r,B} = \{\sigma \in \omega_r : \overline{U_F^\sigma} \cap B \neq \emptyset\}$. A simple geometric argument shows that $\omega_{r,B}$ is finite: since $\{\overline{U_F^\sigma}\}_{\sigma \in \omega_r}$ is a collection of disjoint sets with diameter at least $((\min_{1 \leq i \leq N} L_i) \cdot a \cdot r)$, only finitely many of them can fit inside the ball centered at the center B, with radius $r + 2br$ (call it $(1 + 2b)B$). But if $\overline{U_F^\sigma} \cap B \neq \emptyset$, any two points in $\overline{U_F^\sigma}$ are at most $2br$ away and there is one point that is no more than r away from the center of B; so $\overline{U_F^\sigma} \subset (1 + 2b)B$. Let $k = \#(\omega_{r,B})$.

We see that

$$\lambda^*(B) = \lambda(\{\tau \in [N]^{\mathbb{N}} : x_\tau \in K_F \cap B\}) \leq \lambda(\{[N]_\sigma : \sigma \in \omega_{r,B}\})$$

because if $x_\tau \in K_F \cap B$, then there exists $\sigma \prec \tau$ such that $x_\tau \in \overline{U_F^\sigma}$; we can further specify that $\sigma \in \omega_{r,B}$. This means

$$\lambda^*(B) \leq \sum_{\sigma \in \omega_{r,B}} \lambda([N]_\sigma) = \sum_{\sigma \in \omega_{r,B}} (L_{n_1} \dots L_{n_q})^s \leq \sum_{\sigma \in \omega_{r,B}} r^s = kr^s.$$

Since r and B were chosen arbitrarily it holds for any set V with $\operatorname{diam}(V) \leq 1$ that $\lambda^*(V) \leq k \cdot \operatorname{diam}(V)^s$.

Now consider any countable cover $\{X_i\}_{i\in\mathbb{N}}$ of K_F with non-empty sets of diameter less than $\varepsilon = 1$. It holds that

$$1 = \lambda^*(K_F) \leq \sum_{i=1}^{\infty} \lambda^*(X_i) \leq k \cdot \sum_{i=1}^{\infty} \operatorname{diam}(X_i)^s.$$

Taking the infimum over all such covers gives $\mathscr{H}_\varepsilon^s(K_F) \geq \frac{1}{k}$. As $\varepsilon \to 0^+$, we get $\mathscr{H}^s(K_F) \geq \frac{1}{k} > 0$.

Combining Parts I and II we obtain $0 < \mathscr{H}^s(K_F) < \infty$, so $\dim_{\mathscr{H}}(K_F) = s$. □

4. Measuring Random Fractals: The General Case

This section contains the primary result of this paper: Theorem 4.3. Note that throughout this section, we use the convention that $0^0 = 0$.

Definition 4.1. Let (F, μ) be a stochastic IFS on $K \subset \mathbb{R}^d$. Then, define the random variables $R_\sigma = \operatorname{Lip}(F_\sigma)$ to be the *r-Lipschitz constants* of (F, μ).

(Here, 'r' is simply short for 'random'.)

The distribution of R_σ is induced by μ in the natural way. The following facts are immediately checked from the definition:

Proposition 4.2. $\mathbb{E}[R_\sigma] < \infty$ *(that is, R_σ is integrable). Additionally,*

(1) R_σ is L^t-bounded for all $t \geq 0$ (that is, $\mathbb{E}[R_\sigma^t] < \infty$ for all $t \geq 0$).
(2) $\{R_{\sigma\bullet 1}, \dots, R_{\sigma\bullet N}\}_{\sigma\in[N]^*}$ is a collection of independent and identically distributed N-tuples of Lipschitz factors with the same distribution as the random N-tuple $\{R_1, \dots, R_N\}$.

This second part of this proposition implies that the random variables $\{R_1, \dots, R_N\}$ effectively characterize a stochastic IFS.

Theorem 4.3. *Let (F, μ) be a stochastic IFS on $K \subset \mathbb{R}^d$ with r-Lipschitz constants $\{R_i\}_{i=1}^N$ such that $\mathbb{E}[R_1^0 + \dots + R_N^0] > 1$. Then $\mathbb{P}(\{K_F : \dim_{\mathscr{H}}(K_F) = s\}) = 1$, where s be the unique real number satisfying*

$$\mathbb{E}[R_1^s + \dots + R_N^s] = 1.$$

Remark 4.4. This is an elegant and surprising parallel between self-similarity and stochastic self-similarity: despite randomness being introduced at each stage of the fractal construction, the Hausdorff dimension of the resultant set is the same almost surely (with probability 1), and can often be evaluated explicitly.

We break Theorem 4.3 into several smaller theorems and work our way towards the main result. Unless mentioned otherwise, all the following results only hold under the same assumptions as Theorem 4.3.

Proposition 4.5. *There exists a unique $s \in [0, d]$ such that*

$$\mathbb{E}[R_1^s + \cdots + R_N^s] = 1.$$

Proof. Consider the continuous function

$$\phi(t) = \mathbb{E}[R_1^t + \cdots + R_N^t].$$

By assumption, $\phi(0) > 1$; and $\phi(d) \leq 1$ because the random sets $F_1(K), \ldots, F_N(K)$ are disjoint, with diameters $R_1^d \cdot \operatorname{diam}(K), \ldots, R_N^d \cdot \operatorname{diam}(K)$, and the sum of these diameters can be no larger than $\operatorname{diam}(K)$. The conclusion follows from the intermediate value theorem. Uniqueness is clear from the fact that ϕ is strictly decreasing. □

Theorem 4.6. *K_F is almost surely a set with Hausdorff dimension less than or equal to s. That is,*

$$\mathbb{P}(\{K_F : \dim_{\mathscr{H}}(K_F) \leq s\}) = 1.$$

Proof. For all and $\sigma \in [N]^*$, define $D_\sigma = \operatorname{diam}(K_F^\sigma)$, where K_F^σ is the random set $F_{\sigma|1} \circ \cdots \circ F_{\sigma|(|\sigma|)}(K)$. Clearly,

$$D_\sigma = \operatorname{diam}(K) \cdot \prod_{i=1}^{|\sigma|} R_{\sigma|i}.$$

For all $n \geq 1$, $t \geq 0$, define

$$M_{t,n} = \sum_{|\sigma|=n} D_\sigma^t.$$

We claim that the sequence $\{M_{s,n}\}_{n=1}^\infty$ is a martingale with respect to the filtration $\mathscr{F}_n = \{F_\sigma : |\sigma| \leq n\}$. Indeed,

$$\begin{aligned}
\mathbb{E}[M_{s,n+1} \mid \mathscr{F}_n] &= \mathbb{E}\left[\sum_{|\sigma|=n+1} D_\sigma^s \mid \mathscr{F}_n\right] \\
&= \mathbb{E}\left[\sum_{|\sigma|=n} D_{\sigma\bullet 1}^s + \cdots + D_{\sigma\bullet N}^s \mid \mathscr{F}_n\right] \\
&= \mathbb{E}\left[\sum_{|\sigma|=n} D_\sigma^s \cdot \left(\sum_{i=1}^N R_i^s\right) \mid \mathscr{F}_n\right] \\
&= \sum_{|\sigma|=n} D_\sigma^s \cdot \mathbb{E}[R_1^s + \cdots + R_N^s] \\
&= M_{s,n} \cdot \mathbb{E}[R_1^s + \cdots + R_N^s] = M_{s,n}.
\end{aligned}$$

By the Martingale Convergence Theorem, $M_s \equiv \lim_{n\to\infty} M_{s,n}$ exists and is bounded for ν-a.e. GFS (and thus, $\mathbb{P}$-a.e. GFS attractor K_F). For all n, $\{K_F^\sigma\}_{|\sigma|=n}$ is a cover of K_F, so by the definition of Hausdorff dimension,

$$\mathscr{H}^s(K_F) \leq \lim_{n\to\infty} \sum_{|\sigma|=n} \operatorname{diam}(K_F^\sigma)^s = \lim_{n\to\infty} M_{s,n} = M_s < \infty$$

for $\mathbb{P}$-a.e. random set K_F. So, the Hausdorff dimension is almost surely no greater than s. □

We now construct the proof that the Hausdorff dimension is at least s. The following lemma regarding M_s comes in handy:

Lemma 4.7. *M_s is an L^p-bounded random variable for all $p > 0$. That is, all moments of M_s are finite.*

Proof. It suffices to show by Fatou's Lemma that for all $p \geq 1$, the sequence $\{M_{s,n}^p\}_{n=1}^\infty$ is uniformly bounded in expectation. Of course, by monotonicity of L^p norms, it suffices to show this for all integers $p \geq 2$.

Begin with the $p = 2$ case. For $i = 1, \ldots, N$, define $M_{s,n}^{(i)}$ to be an iid copy of $M_{s,n}$. Then, by shifting our entire construction down one step, we have $M_{s,n+1} = \sum_{i=1}^N R_i^s M_{s,n}^{(i)}$. Clearly $M_{s,1}$ has finite moments of all integer orders:

$$\begin{aligned}
\mathbb{E}[M_{s,1}] &= \mathbb{E}\left[\sum_{i=1}^N D_{(i)}^s\right] = \operatorname{diam}(K) \cdot \mathbb{E}\left[R_1^s + \cdots + R_N^s\right] = \operatorname{diam}(K). \\
\mathbb{E}[M_{s,1}^p] &= \mathbb{E}\left[\left(\sum_{i=1}^N R_i^s \cdot \operatorname{diam}(K)\right)^p\right] = \operatorname{diam}(K)^p \cdot \mathbb{E}\left[\left(\sum_{i=1}^N R_i^s\right)^p\right] \\
&\leq \operatorname{diam}(K)^p \cdot N^{sp}.
\end{aligned}$$

Now, for $n \geq 1$ we observe that

$$\begin{aligned}
\mathbb{E}[M_{s,n+1}^2] &= \mathbb{E}\left[\left(\sum_{i=1}^N R_i^s M_{s,n}^{(i)}\right)^2\right] \\
&= \mathbb{E}\left[\sum_{i=1}^N R_i^{2s} M_{s,n}^2\right] + \mathbb{E}\left[\sum_{i \neq j} R_i^s R_j^s M_{s,n}^{(i)} M_{s,n}^{(j)}\right],
\end{aligned}$$

which we obtain by splitting the sum into the diagonal and non-diagonal terms. Here, $M_{s,n}^{(i)}$ is independent of R_j for all i, j; and so,

$$\mathbb{E}[M_{s,n+1}^2] = \mathbb{E}\left[\sum_{i=1}^N R_i^{2s}\right] \mathbb{E}\left[M_{s,n}^2\right] + \operatorname{diam}(K)^2 \cdot \mathbb{E}\left[\sum_{i \neq j} R_i^s R_j^s\right].$$

The diagonal term at the end can be split up as

$$\mathbb{E}\left[\sum_{i \neq j} R_i^s R_j^s\right] = \mathbb{E}\left[\left(\sum_{i=1}^N R_i^s\right)^2\right] - \mathbb{E}\left[\sum_{i=1}^N R_i^{2s}\right] \equiv C_2 - c_2.$$

Then, letting $a_n = \mathbb{E}[M_{s,n}^2]$ we may note that

$$a_{n+1} = c_2 a_n + \operatorname{diam}(K)^2 (C_2 - c_2).$$

This is a fairly simple linear recurrence relation, and solving down to $n = 1$, we get

$$\begin{aligned}
a_n &= c_2 a_{n-1} + \operatorname{diam}(K)^2 (C_2 - c_2) \\
&= c_2^2 a_{n-2} + c_2(\operatorname{diam}(K)^2 (C_2 - c_2)) + \operatorname{diam}(K)^2 (C_2 - c_2) \\
&= c_2^3 a_{n-3} + \operatorname{diam}(K)^2 (C_2 - c_2)(c_2^2 + c_2 + 1) = \ldots \\
&= c_2^{n-1} a_1 + \operatorname{diam}(K)^2 (C_2 - c_2) \sum_{j=0}^{n-2} c_2^j.
\end{aligned}$$

As shown above, $a_1 = \mathbb{E}[M_{s,1}^2] \leq \operatorname{diam}(K)^2 \cdot N^{2s}$. Furthermore, $c_2 < 1$ because the function $\phi(t)$ mentioned in Proposition 4.5 is strictly decreasing and $s < 2s$. Thus, for all n,

$$\begin{aligned}\mathbb{E}[M_{s,n}^2] = a_n &\leq c_2^{n-1}\operatorname{diam}(K)^2 \cdot N^{2s} + \operatorname{diam}(K)^2(C_2 - c_2) \cdot \frac{1}{1-c_2} \\ &\leq \operatorname{diam}(K)^2 \cdot N^{2s} + \operatorname{diam}(K)^2(C_2 - c_2) \cdot \frac{1}{1-c_2} < \infty,\end{aligned}$$

so $\{M_{s,n}\}$ is an L^2-bounded sequence.

For the $p \geq 2$ case, we outline the strong induction method that proves finiteness of moments. Suppose p is an integer such that $\{M_{s,n}\}$ is L^k-bounded for all integers $k = p-1, p-2, \ldots, 1$. Then,

$$\mathbb{E}[M_{s,n+1}^p] = \mathbb{E}\left[\left(\sum_{i=1}^{N} R_i^s \cdot M_{s,n}^{(i)}\right)^p\right]$$

as before. We again split into all the diagonal and non-diagonal terms. There are exactly $N^p - N$ non-diagonal terms, and each one is a product of lower-order moments of $M_{s,n}$, which by the induction hypothesis are all finite; so the product itself, by Hölder's inequality, is finite. Thus the non-diagonal term is some finite constant B_p. The diagonal terms sum to

$$\mathbb{E}\left[\sum_{i=1}^{N} R_i^{sp} M_{s,n}^p\right] = \mathbb{E}\left[\sum_{i=1}^{N} R_i^{sp}\right] \cdot \mathbb{E}[M_{s,n}^p] \equiv c_k \mathbb{E}[M_{s,n}^p]$$

by independence. The expected sum of all R_i^{sp} is less than 1 since the function ϕ from Proposition 4.5 is strictly decreasing; so we again have a linear recurrence relation

$$\mathbb{E}[M_{s,n+1}^p] = c_p \mathbb{E}[M_{s,n}^p] + B_p.$$

Since B_p is finite, $c_p < 1$, and $\mathbb{E}[M_{s,1}^p] < \infty$, we again calculate a uniform bound on $\{M_{s,n}^p\}$ as with the $p = 2$ case. Thus, $\{M_{s,n}\}$ is an L^p-bounded sequence for all $p > 0$, and so M_s has finite moments of all orders. □

Now we introduce a random variable taking values in $C_c(\mathbb{R}^d)^*$, the dual of the space of continuous functions on $\mathbb{R}^d$ with compact support. For $f \in C_c(\mathbb{R}^d)$, define

$$G(f) = \lim_{n\to\infty} \sum_{|\sigma|=n} f(x_\sigma) D_\sigma^s,$$

where x_σ is some point in K_F^σ. $G(f)$ is well-defined, because any infinite nested sequence of sets $(K_F^{\sigma_i})_{i=1}^\infty$ can be associated with the sequence $\tau \in [N]^{\mathbb{N}}$ such that $\tau = \bigcap_i [N]_{\sigma_i}$. Consequently, $x_{\sigma_i} \to x_\tau$ and so $f(x_{\sigma_i}) \to f(x_\tau)$.

Theorem 4.8. *For ν-a.e. GFS F and all $f \in C_c(K)$, $G(f) \in C_c(K)^*$ and the norm of F is given by $\|G\| = M_s$.*

Proof. For each $\sigma \in [N]^*$, define intermediary random variables

$$M_{\sigma,n} = \sum_{|\omega|=n} \prod_{q=1}^{n} R_{\sigma\bullet[\omega|q]}^s \quad \text{and} \quad M_\sigma = \lim_{n\to\infty} M_{\sigma,n}.$$

$M_{\sigma,n}$ can be thought of as the analog of $M_{s,n}$ if we perform the fractal generating process, starting at node σ and normalizing the diameter of the initial set to 1.

This makes it easy to see that $M_{\sigma,n}$ is equal in distribution to $\dfrac{M_{s,n}}{\text{diam}(K)^s}$; so M_σ is equal in distribution to $\dfrac{M_s}{\text{diam}(K)^s}$ by the continuous mapping theorem on random variables. This also means $\{M_\sigma\}_{\sigma\in[N]^*}$ is a mutually independent family, itself independent of $\mathscr{F}_n$. Finally, for every fixed n and q,

$$M_{s,n+q} = \sum_{\sigma\in[N]^n} D_\sigma^s M_{\sigma,q} \overset{q\to\infty}{\Longrightarrow} M_s = \sum_{\sigma\in[N]^n} D_\sigma^s \cdot M_\sigma.$$

Now say $f \in C_c(K)$, and for $p, r \in \mathbb{N}$, define

$$\varepsilon_{p,r} = \left| \sum_{|\sigma|=p} f(x_\sigma) D_\sigma^s - \sum_{|\sigma|=r} f(x_\sigma) D_\sigma^s \right|.$$

Suppose $p, r > k$ for some $k \in \mathbb{N}$; then,

$$\begin{aligned}
\varepsilon_{p,r} &= \left| \sum_{|\sigma|=k} D_\sigma^s \left(\sum_{|\omega|=p-k} f(x_{\sigma\bullet\omega}) \prod_{q=1}^{p-k} R_{\sigma\bullet[\omega|q]} - \sum_{|\omega|=r-k} f(x_{\sigma\bullet\omega}) \prod_{q=1}^{r-k} R_{\sigma\bullet[\omega|q]} \right) \right| \\
&\leq \sum_{|\sigma|=k} D_\sigma^s \Bigg(\sup_{|\omega|=p-k} |f(x_{\sigma\bullet\omega}) - f(x_\sigma)| M_{\sigma,p-k} + \sup_{|\omega|=r-k} |f(x_{\sigma\bullet\omega}) - f(x_\sigma)| M_{\sigma,r-k} \cdot \\
&\quad |f(x_\sigma)| \cdot |M_{\sigma,p-k} - M_{\sigma,r-k}| \Bigg) \\
&\leq \sum_{|\sigma|=k} D_\sigma^s \cdot (\text{diam}(f(K_F^\sigma))(M_{\sigma,p-k} + M_{\sigma,r-k})) + \|f\| \cdot |M_{\sigma,p-k} - M_{\sigma,r-k}|).
\end{aligned}$$

As $k, p, r \to \infty$, $\text{diam}(K_F^\sigma) \to 0$ for ν-a.e. GFS (since one of the Lipschitz factors is almost surely less than 1) and $M_{\sigma,p-k}, M_{\sigma,r-k} \to M_\sigma$. Thus, $\varepsilon_{p,r} \to 0$ almost surely.

It is immediate that G is linear; and whenever $K \subset f^{-1}(\{1\})$,

$$G(f) = \lim_{n\to\infty} \sum_{|\sigma|=n} f(x_\sigma) D_\sigma^s = \lim_{n\to\infty} \sum_{|\sigma|=n} D_\sigma^s = \lim_{n\to\infty} M_{n,s} = M_s.$$

Whenever $\|f\| \leq 1$, $\|G(f)\| \leq M_s$; and as shown, the value M_s is in fact attained, so $\|G\| = M_s$. □

By the Riesz-Markov theorem, we may associate G with a (random) measure γ on $\mathbb{R}^d$ that satisfies $G(f) = \int f \, d\gamma$ for all f. γ can intuitively be understood as a stochastic counterpart to the s-dimensional Hausdorff measure. We will carefully study the properties of γ for the next few theorems.

Theorem 4.9. *If $A \in \mathscr{K}(\mathbb{R}^d)$, then*

$$\sum_{|\sigma|=n, K_F^\sigma \cap A \neq \emptyset} D_\sigma^s M_\sigma \downarrow \gamma(A) \quad \textit{as } n \to \infty$$

for ν-a.e. GFS F.

Proof. Say $k \in \mathbb{N}$ and $\varepsilon > 0$. Almost surely $\sum_{|\sigma|=k} D_\sigma^s M_\sigma$ is finite, so there exists a set $\Theta \subsetneq [N]^k$ such that $\sum_{\sigma\in[N]^k\setminus\Theta} D_\sigma^s M_\sigma < \varepsilon$. By Urysohn's Lemma we can find $f \in C_c(\mathbb{R}^d)$ such that $f^{-1}(\{1\}) = A$ and $K_F^\sigma \subset f^{-1}(\{0\})$ whenever $\sigma \in \Theta$ and $K_F^\sigma \cap A = \emptyset$. Then,

$$\begin{aligned}
\gamma(A) &\leq \int f \; d\gamma = G(f) = \lim_{n\to\infty} \sum_{|\sigma|=n} f(x_\sigma) D_\sigma^s \\
&= \lim_{n\to\infty} \sum_{|\sigma|=k} \sum_{|\omega|=n-k} f(x_{\sigma\bullet\omega}) D_{\sigma\bullet\omega}^s \\
&\leq \lim_{n\to\infty} \sum_{|\sigma|=k, K_F^\sigma \cap A \neq \emptyset} \sum_{|\omega|=n-k} 1 \cdot D_{\sigma\bullet\omega}^s + \lim_{n\to\infty} \sum_{\sigma \in [N]^k \setminus \Theta, K_F^\sigma \cap A = \emptyset} \sum_{|\omega|=n-k} 1 \cdot D_{\sigma\bullet\omega}^s \\
&= \lim_{n\to\infty} \sum_{|\sigma|=k, K_F^\sigma \cap A \neq \emptyset} D_\sigma^s M_{\sigma,n-k} + \lim_{n\to\infty} \sum_{\sigma \in [N]^k \setminus \Theta, K_F^\sigma \cap A = \emptyset} D_\sigma^s M_{\sigma,n-k}.
\end{aligned}$$

The second term is at most ε; and the first term, as $n \to \infty$, is $\sum_{|\sigma|=k, K_F^\sigma \cap A \neq \emptyset} D_\sigma^s M_\sigma$. Thus,

$$\gamma(A) \leq \sum_{|\sigma|=k, K_F^\sigma \cap A \neq \emptyset} D_\sigma^s M_\sigma.$$

Now, note that

$$\sum_{|\sigma|=k+1, K_F^\sigma \cap A \neq \emptyset} D_\sigma^s M_\sigma \leq \sum_{|\sigma|=k, K_F^\sigma \cap A \neq \emptyset} D_\sigma^s M_\sigma,$$

since $K_F^{\sigma\bullet i} \subset K_F^\sigma$ for all i; thus, the sequence in the theorem is decreasing and bounded below uniformly by $\gamma(A)$, and its limit exists. Define the measure $\tilde{\gamma}(A)$ as this limit. Now say $B \in \mathscr{K}(\mathbb{R}^d)$ and $A \cap B = \emptyset$; then, for large enough k, $|\sigma| = k$ implies that K_F^σ cannot intersect both A and B; so $\tilde{\gamma}(A) + \tilde{\gamma}(B) \leq M_s$.

Now, take an increasing sequence of compact sets $\{B_n\}_{n=1}^\infty$ such that $\gamma(B_n) \to \gamma(\mathbb{R}^d \setminus A)$. For each n, we have

$$\gamma(A) + \gamma(B_n) \leq \tilde{\gamma}(A) + \gamma(B_n) \leq M_s.$$

As $n \to \infty$, the left-hand side of this equation becomes $\gamma(A) + \gamma(\mathbb{R}^d \setminus A) = M_s$, since $\gamma(\mathbb{R}^d) = M_s$. So we in fact have

$$M_s \leq \tilde{\gamma}(A) + \gamma(\mathbb{R}^d \setminus A) \leq M_s.$$

This implies that for all compact sets A,

$$\gamma(A) = \tilde{\gamma}(A) = \lim_{n\to\infty} \sum_{|\sigma|=n, K_F^\sigma \cap A \neq \emptyset} D_\sigma^s M_\sigma,$$

so we're done. □

Theorem 4.10. *For $\mathbb{P}$-a.e. random set K_F,*

$$\gamma(K_F) = \gamma(\mathbb{R}^d) = M_s.$$

Proof. Clearly,

$$\gamma(\mathbb{R}^d) = \int_{\mathbb{R}^d} 1 \; d\gamma = G(1) = \lim_{n\to\infty} \sum_{|\sigma|=n} D_\sigma^s = \lim_{n\to\infty} M_{s,n} = M_s.$$

Now, to find the measure of K_F, we may apply the standard result about measures of decreasing sequences of sets on finite measure spaces (since M_s is almost surely finite):

$$\gamma(K_F) = \lim_{q\to\infty} \gamma\left(\bigcup_{|\sigma|=q} K_F^\sigma\right).$$

Fixing q, the random set $K_q = \bigcup_{|\sigma|=q} K_F^\sigma$ is compact, so by Theorem 4.9,

$$\gamma(K_q) = \lim_{n\to\infty} \sum_{|\sigma|=n, K_q\cap K_F^\sigma\neq\emptyset} D_\sigma^s M_\sigma.$$

But for large enough n, $K_q \supset K_n \supset K_F^\sigma$ for all σ with $|\sigma| = n$; so the limit on the right-hand side is simply $\lim_{n\to\infty} \sum_{|\sigma|=n} D_\sigma^s M_\sigma = M_s$. So, $\gamma(K_q) = M_s$. $\square$

Note that because of how we defined a contraction, we run into the issue that F_σ may be a constant map with $R_\sigma = 0$. To address this case we introduce a variable $M_0 = \lim_{n\to\infty} M_{0,n} = \lim_{n\to\infty} \sum_{|\sigma|=n} D_\sigma^0$. The following theorem provides a useful characterization of when the resultant set K_F is empty.

$M_{0,n}$ counts the number of non-empty sets at the n-th stage of the random fractal generation; and what this theorem says is that $M_{0,n}$ must grow without bound for the resultant set to be non-empty.

Theorem 4.11. *For ν-a.e. GFS, $K_F \neq \emptyset$ if and only if $M_0 = \infty$.*

Proof. We first show that $M_0 = 0$ or $M_0 = \infty$ almost surely. This follows from the fact that $\{M_{0,n}\}$ is a Galton-Watson branching process, wherein each node has at most N children. From μ, we can obtain a probability vector $(p_0^\sigma, \dots, p_N^\sigma)$, where p_i^σ represents the probability $\mathrm{Lip}(F_{\sigma\bullet j}) > 0$ for exactly i indices. Of course, by our construction, the collection $\{p_0^\sigma, \dots, p_N^\sigma\}_{\sigma\in[N]^*}$ is iid.

Define $f(x) = \sum_{k=0}^N p_k^\sigma x^k$. Then, $f'(1) = \sum_{k=0}^N kp_k^\sigma = \mathbb{E}[\#\text{ of surviving nodes}]$. But this is exactly $\mathbb{E}[R_1^0 + \cdots + R_N^0]$, which is greater than 1 by hypothesis. Thus, $\{M_{0,n}\}$ is supercritical; by a standard result to do with branching processes, $\lim_{n\to\infty} M_{0,n} \in \{0, \infty\}$ almost surely.

But now the theorem is fairly obvious. If $M_0 = \infty$, it is clear that $K_F \neq \emptyset$. If $M_0 < \infty$, $M_0 = 0 \implies K_F = \emptyset$. Thus we are done. $\square$

Note that this theorem also explains why $\mathbb{E}[R_1^0 + \cdots + R_N^0] > 1$ in our hypothesis. If this quantity was less than 1, the Galton-Watson process would reach extinction almost surely; the same would happen if $\mathbb{E}[R_1^0 + \cdots + R_N^0] = 1$ and $p_0^\sigma \neq 0$.

Theorem 4.12. *If $\mathbb{E}[M_s] > 0$, then*

$$\mathbb{P}(\{K_F : \gamma(K_F) > 0 \text{ and } K_F \neq \emptyset\}) = \mathbb{P}(\{K_F : K_F \neq \emptyset\}).$$

In other words, if K_F is non-empty, its γ-measure is strictly positive almost surely.

Proof. Of course, if $\mathbb{E}[M_s] > 0$, M_s is strictly positive with probability $\delta > 0$. Additionally, $\sum_{i=1}^N R_i^0 \leq N$, so the set $X_n = \{|\sigma| = n : D_\sigma > 0\}$ is finite.

Suppose F is some arbitrarily subset of $[N]^n$. Then,

$$P(M_s = 0 \text{ and } F = X_n) = P(M_\sigma = 0 \text{ for } \sigma \in F \text{ and } F = X_n).$$

This holds because if D_σ and M_σ are both strictly positive, $M_s > 0$. By independence of the family $\{M_\sigma\}_{\sigma\in[N^*]}$, the above probability works out to be equal to

$$P(M_\sigma = 0)^{\#F} \cdot P(F = X_n) = (1-\delta)^{\#F} P(F = X_n),$$

since M_σ has the same distribution as $M_s/\operatorname{diam}(K)^s$. Now, for fixed q, we obtain

$$\begin{aligned} P(M_s = 0 \text{ and } \#X_n \geq q) &= \sum_{\#F \geq q} P(M_s = 0 \text{ and } F = X_n) \\ &= \sum_{\#F \geq q} (1-\delta)^{\#F} P(F = X_n) \\ &\leq (1-\delta)^q \cdot P(\#X_n \geq q) \\ \implies P(M_s > 0 \text{ and } \#X_n \geq q) &\geq [1-(1-\delta)^q] \cdot P(\#X_n \geq q). \end{aligned}$$

But note that $\#X_n = M_{0,n}$ as introduced in 4.6, so we really have the inequality

$$P(M_s > 0) \geq P(M_s > 0 \text{ and } M_{0,n} \geq q) \geq [1-(1-\delta)^q] \cdot P(M_{0,n} \geq q).$$

Taking $n \to \infty$ we get

$$P(M_s > 0) \geq [1-(1-\delta)^q] P(M_0 \geq q).$$

Then, taking $q \to \infty$ (which we can now do since $q \leq N^n$; N is fixed but $n \to \infty$), we get

$$P(M_s > 0) \geq P(M_0 = \infty).$$

Now suppose $M_0 = 0$. Since $M_{0,n}$ can only take on integer values, it must hold that for large enough n, $M_{0,n} = 0$. But then $D_\sigma = 0$ for all $|\sigma| = n$, and so $M_s = 0$; so $M_s > 0 \implies M_0 > 0 \implies M_0 = \infty$ by Theorem 4.11.

So, $P(M_s > 0) \geq P(M_0 = \infty)$, but $M_s > 0$ only if $M_0 = \infty$; this implies that $M_s > 0$ iff $M_0 = \infty$. But $M_0 = \infty$ iff $K_F \neq \emptyset$, and $M_s > 0$ iff $\gamma(K_F) > 0$; so we have $K_F \neq \emptyset \iff \gamma(K_F) > 0$, meaning we're done. $\square$

Theorem 4.13. *In addition to the assumptions of Theorem 4.3, assume there exists $C > 0$ such that $P(R_\sigma > C \mid R_\sigma > 0) = 1$. Then, if $t < s$ and $E \subset \mathbb{R}^d$ is compact, $\mathscr{H}^t(E) < \infty \implies \gamma(E) = 0$.*

Proof. By Lemma 4.7, M_s is L^p-bounded for all p; since M_σ has the same distribution as $M_s/\operatorname{diam}(K)^s$, the same holds for M_σ. Now, fixing some $k > 0$ and $t < s$, Chebyshev's inequality with any $r > 0$ gives

$$P(D_\sigma^s M_\sigma > k D_\sigma^t) = P(D_\sigma^{s-t} M_\sigma > k) \leq \frac{\mathbb{E}[(D_\sigma^{s-t} M_\sigma)^r]}{k^r} = \frac{\mathbb{E}[D_\sigma^{r(s-t)}]\mathbb{E}[M_\sigma^r]}{k^r},$$

the last equality holding since D_σ and M_σ are independent. This means

$$\begin{aligned} \sum_{|\sigma|=n} P(D_\sigma^s M_s > k D_\sigma^t) &\leq \sum_{|\sigma|=n} \frac{\mathbb{E}[D_\sigma^{r(s-t)}]\mathbb{E}[M_\sigma^r]}{k^r} \\ &= \sum_{|\sigma|=n} \frac{\mathbb{E}[D_\sigma^{r(s-t)}]\mathbb{E}[M_s^r]}{\operatorname{diam}(K)^{sr} k^r} \\ &= \frac{\mathbb{E}[\sum_{|\sigma|=n} D_\sigma^{r(s-t)}]\mathbb{E}[M_s^r]}{\operatorname{diam}(K)^{sr} k^r} \\ &= \frac{\mathbb{E}[M_{r(s-t),n}]\mathbb{E}[M_s^r]}{\operatorname{diam}(K)^{sr} k^r} \\ \implies P(\exists \sigma \text{ s.t. } |\sigma| = n, D_\sigma^s M_s > k D_\sigma^t) &\leq \frac{\mathbb{E}[M_{r(s-t),n}]\mathbb{E}[M_s^r]}{\operatorname{diam}(K)^{sr} k^r}. \end{aligned}$$

The argument from Theorem 4.6 shows that

$$\mathbb{E}[M_{r(s-t),n}] = \mathbb{E}[M_{r(s-t),n-1}] \cdot \mathbb{E}[R_1^{r(s-t)} + \cdots + R_N^{r(s-t)}].$$

Iterating downwards,

$$\mathbb{E}[M_{r(s-t),n}] = \mathbb{E}[M_{r(s-t),1}] \cdot \mathbb{E}\left[\sum_{i=1}^{N} R_i^{r(s-t)}\right]^{n-1} = \mathbb{E}\left[\sum_{i=1}^{N} R_i^{r(s-t)}\right]^{n}.$$

Choose r large enough that $\mathbb{E}[\sum_{i\leq N} R_i^{r(s-t)}] < 1$ (here it suffices that $r(s-t) > s$); then,

$$\begin{aligned}\sum_{n=1}^{\infty} P(\exists\sigma \text{ s.t. } |\sigma| = n, D_\sigma^s M_s > kD_\sigma^t) &\leq \sum_{n=1}^{\infty} \frac{\mathbb{E}[M_{r(s-t),n}]\mathbb{E}[M_s^r]}{\operatorname{diam}(K)^{sr}k^r} \\ &= \sum_{n=1}^{\infty} \frac{\mathbb{E}[\sum_{i\leq N} R_i^{r(s-t)}]^n \mathbb{E}[M_s^r]}{\operatorname{diam}(K)^{sr}k^r} \\ &= \frac{\mathbb{E}[M_s^r]}{\operatorname{diam}(K)^{sr}k^r} \sum_{i=1}^{\infty} \mathbb{E}\left[\sum_{i=1}^{N} R_i^{r(s-t)}\right]^n < \infty.\end{aligned}$$

By the Borel-Cantelli lemma,

$$P(\exists \text{infinitely many } n \text{ s.t. for some } |\sigma| = n, D_\sigma^s M_\sigma > kD_\sigma^t) = 0.$$

This implies that almost surely, there are only finitely many such n; and thus a maximum such n. So, there exists Q such that $|\sigma| \geq Q \implies D_\sigma^s M_\sigma \leq kD_\sigma^t$. Define $\varepsilon = \min\{1, \min\{D_\sigma > 0 : \sigma \in [N]^Q\}\}$.

Let K_F be a particular realization of the random set generated by this fractal generating process. For all $\tau \in [N]^{\mathbb{N}}$ and $p \in \mathbb{N}$, define

$$\Theta_{\tau,p} = \{\sigma \in [N^*] : \sigma \prec \tau, D_\sigma < 2^{-p} \leq D_{\sigma|(|\sigma|-1)}\}.$$

Let $E \subset \mathbb{R}^d$ be compact so that $\mathscr{H}^t(E) < \infty$. Choose a sequence $\{S_n\}$ of d-dimensional spheres such that $E \subset \bigcup_n S_n$ and $\operatorname{diam}(S_n) < \frac{\varepsilon}{2}$ for all n. For every n, there exists $p_n \in \mathbb{N}$ such that $2^{-1-p_n} \leq \operatorname{diam}(S_n) < 2^{-p_n}$; so define, for every n, the set

$$A_n = \bigcup\{\Theta_{\tau,p_n} : x_\tau \in S_n \cap E\}.$$

Now suppose σ_1 are σ_2 two distinct elements of A_n, and say there exists η such that $\sigma_1 \bullet \eta = \sigma_2$. Then, there exists $\tau \in [N]^{\mathbb{N}}$ such that $\sigma_2 \prec \tau$, meaning $\sigma_1 \prec \tau$; so, $\sigma_1, \sigma_2 \in \Theta_{\tau,p_n}$ for some τ. But this would imply

$$D_{\sigma_2} = D_{\sigma_1\bullet\eta} < 2^{-p_n} \leq D_{\sigma_1\bullet[\eta|(|\eta|-1)]} < \cdots < D_{\sigma_1} < 2^{-p_n},$$

an absurdity. Thus there are no two such elements in A_n. What this implies is that if $\sigma_1, \sigma_2 \in A_n$ and $\sigma_1 \neq \sigma_2$, $\operatorname{int}(K_F^{\sigma_1})$ and $\operatorname{int}(K_F^{\sigma_2})$ are disjoint. For any given $x \in S_n$ and $y \in K_F^\sigma$ where $\sigma \in A_n$, we have $\sigma \in \Theta_{\tau,p_n}$ for some τ such that $x_\tau \in S_n$. x_τ and x are at most $\operatorname{diam}(S_n) < 2^{-p_n}$ away; and since $x_\tau \in K_F^\sigma$, x_τ and y are no more than $\operatorname{diam}(K_F^\sigma) < 2^{-p_n}$ away. Thus, $d(x, y) < 2^{1-p_n}$. This implies that for any $x \in S_n$,

$$\bigcup_{\sigma\in A_n} K_F^\sigma \subset B(x, 2^{1-p_n}).$$

Note, however, that if $\sigma \in \Theta_{\tau,p_n}$, then $\operatorname{diam}(K_F^\sigma) \geq C \cdot \operatorname{diam}(K_F^{\sigma|(|\sigma|-1)}) \geq C \cdot 2^{-p_n}$, since $R_\sigma > C$ almost surely whenever $R_\sigma > 0$. Letting $\mathscr{L}^d$ denote the d-dimensional Lebesgue measure, we therefore obtain

$$\mathscr{L}^d(\operatorname{int}(K_F^\sigma)) \geq \mathscr{L}^d(\operatorname{int}(K)) \cdot \left(\frac{\operatorname{diam}(K_F^\sigma)}{\operatorname{diam}(K)}\right)^d \geq \mathscr{L}^d(\operatorname{int}(K)) \cdot \left(\frac{C \cdot 2^{-p_n}}{\operatorname{diam}(K)}\right)^d.$$

However, since the sets $\{K_F^\sigma\}_{\sigma \in A_n}$ have disjoint interiors and all fit inside the ball $B(x, 2^{1-p_n})$, we have

$$\#A_n \cdot \mathscr{L}^d(\operatorname{int}(K)) \cdot \left(\frac{C \cdot 2^{-p_n}}{\operatorname{diam}(K)}\right)^d \leq \mathscr{L}^d(B(x, 2^{1-p_n})) = 2^{-p_n d}\mathscr{L}^d(B(0,2)).$$

This implies that

$$\#A_n \leq \frac{\mathscr{L}^d(B(0,2))}{\mathscr{L}^d(\operatorname{int}(K))} \cdot \left(\frac{\operatorname{diam}(K)}{C}\right)^d \equiv J.$$

For all $\sigma \in A_n$, though,

$$D_\sigma^s M_\sigma \leq k D_\sigma^t < k(2^{-p_n})^t \leq 2^t \operatorname{diam}(S_n)^t k.$$

For each n, let $q(n) = \max\{|\sigma| : \sigma \in A_n\}$ (which exists because A_n is finite). Then,

$$\gamma(E \cap S_n) \leq \sum_{|\sigma|=q(n), E\cap S_n \cap K_F^\sigma \neq \emptyset} D_\sigma^s M_\sigma$$

by Theorem 4.9. For each term in this sum, take the corresponding σ; there exists a $\sigma_* \in A_n$ such that $\sigma \mid (|\sigma_*|) = \sigma_*$, and therefore, $K_F^\sigma \subset K_F^{\sigma_*}$. Furthermore, $D_{\sigma_*}^s M_{\sigma_*} \geq D_\sigma^s M_\sigma$; so, we in fact have

$$\begin{aligned}\gamma(E \cap S_n) &\leq \sum_{\sigma \in A_n} D_\sigma^s M_\sigma \leq \#A_n \cdot k \cdot 2^t \operatorname{diam}(S_n)^t \\ &\leq kJ2^t \operatorname{diam}(S_n)^t \\ \implies \gamma(E) &\leq \sum_{n=1}^\infty kJ2^t \operatorname{diam}(S_n)^t \leq kJ2^t \mathscr{H}^t(E).\end{aligned}$$

Since, at the beginning of the proof, we fixed $k > 0$ arbitrarily, we take $k \to 0$ to get $\gamma(E) = 0$ whenever $\mathscr{H}^t(E) < \infty$. □

We now come to the final theorem that, combined with Theorem 4.6, proves Theorem 4.3.

Theorem 4.14. *Suppose K_F is non-empty. Then, K_F is almost surely a set with Hausdorff dimension greater than or equal to s. That is,*

$$\mathbb{P}(K_F : K_F \neq \emptyset \text{ and } \dim_{\mathscr{H}}(K_F) \geq s) = \mathbb{P}(K_F : \dim_{\mathscr{H}}(K_F) \neq \emptyset).$$

In other words,

$$P(\dim_{\mathscr{H}}(K_F) \geq s \mid K_F \neq \emptyset) = 1.$$

Proof. For each $n \in \mathbb{N}$, we define an auxiliary stochastic IFS (F_n, μ_n) meeting the conditions of Theorem 4.13. Consider the map $\psi_n : \tilde{\mathscr{S}}(\mathbb{R}^d)^N \to \tilde{\mathscr{S}}(\mathbb{R}^d)^N$ defined as follows: $\psi_n(g_1, \dots, g_N) = (g_1^{(n)}, \dots, g_N^{(n)})$, where

$$g_i^{(n)}(x) = \begin{cases} g_i(x) & \text{if } \operatorname{Lip}(g_i) \geq \frac{1}{n}, \\ y & \text{if } \operatorname{Lip}(g_i) < \frac{1}{n}, \end{cases}$$

where y is a randomly chosen element not in K. (Here, we should really be dealing with similitudes on the compact set K, not $\mathbb{R}^d$; but any similitude on K can be extended to one on $\mathbb{R}^d$, so this suffices.)

Then let $\mu_n = \psi_{n\#\mu}$, the pushforward measure of μ under the map ψ_n. Refer to the resultant measure on $\mathscr{K}(\mathbb{R}^d)$ as P_n. For all $i \leq N$, the induced r-Lipschitz constants are $R_i^{(n)} = R_i \cdot \mathbb{1}_{\{R_i \geq 1/n\}}$, where $\mathbb{1}_A$ represents the indicator random variable for an event A.

Now note that $\psi_n(g_1, \ldots, g_N) \to g_1, \ldots, g_N$ as $n \to \infty$ for all N-tuples of contractions; therefore, the measures μ_n converge weakly to μ. To see why, we observe that for any bounded continuous ξ on $\tilde{\mathscr{S}}(\mathbb{R}^d)^N$,

$$\int \xi \, d\mu_n = \int \xi \circ \psi_n \, d\mu \to \int \xi \circ \mathrm{id} \, d\mu = \int \xi \, d\mu.$$

The convergence is a consequence of the dominated convergence theorem, because $|\xi \circ \psi_n| \leq |\xi| \leq \|\xi\|$. By the same reasoning, $\{P_n\}$ converges weakly to P. Lastly, as $n \to \infty$, $R_i^{(n)} = R_i \cdot \mathbb{1}_{\{R_i \geq 1/n\}} \uparrow R_i \cdot \mathbb{1}_{\{R_i > 0\}} = R_i$; so, by the monotone convergence theorem, for all $t > 0$,

$$\mathbb{E}\left[{R_1^{(n)}}^t + \cdots + {R_N^{(n)}}^t\right] \uparrow \mathbb{E}[R_1^t + \cdots + R_N^t].$$

This means there exists some n_0 so large that for all $n \geq n_0$,

$$\mathbb{E}\left[{R_1^{(n)}}^t + \cdots + {R_N^{(n)}}^t\right] > 1.$$

For each n, let $F_{n,\sigma}$ equal $F_n(\sigma)$ be the contraction map at node σ for the nth auxiliary IFS; and let

$$K_{F,n}^\sigma = F_{n,(\sigma|1)} \circ \cdots \circ F_{n,\sigma}(K).$$

Define $D_{n,\sigma} = \mathrm{diam}(K_{F,n}^\sigma)$. It follows that

$$K_{F,n}^\sigma = \begin{cases} K_F^\sigma & \text{if } R_{\sigma|1}^{(n)}, \ldots, R_\sigma^{(n)} \geq 1/n, \\ \{y\} & \text{else,} \end{cases}$$

and consequently,

$$D_{n,\sigma} = \begin{cases} D_\sigma & \text{if } R_{\sigma|1}^{(n)}, \ldots, R_\sigma^{(n)} \geq 1/n, \\ 0 & \text{else.} \end{cases}$$

The resultant compact set will be denoted $K_{F,n} = \bigcap_{q=1}^\infty \bigcup_{|\sigma|=q} K_{F,n}^\sigma \cap K$ (by intersecting with K we no longer have to worry about the singleton $\{y\}$). For all $n \in \mathbb{N}$ and $\sigma \in [N]^*$, $K_{F,n}^\sigma \subset K_{F,n+1}^\sigma \subset K_F^\sigma$ for ν-a.e. GFS F; therefore, for all $n \in \mathbb{N}$, $K_{F,n} \subset K_{F,n+1} \subset K_F$.

For all $n \in \mathbb{N}$ and $p \leq N$, let $q_n = P(K_{F,n} = \emptyset)$, $C_{n,p} = P(\sum_{i \leq N} {R_i^{(n)}}^0 = p)$. Correspondingly, let $q_0 = P(K_F = \emptyset)$ and $C_{0,p} = P(\sum_{i \leq N} R_i^0 = p)$. Fixing n, we get

$$q_n = P(K_{F,n} = \emptyset) = \sum_{p=0}^N P\left(K_{F,n} = \emptyset \text{ and } \sum_{i \leq N} {R_i^{(n)}}^0 = p\right).$$

For each $p \leq N$ (including 0), if $\sum_{i \leq N} {R_i^{(n)}}^0 = p$, there are p non-empty, non-singleton sets after the first N-tuple of contractions is applied. If $K_{F,n} = \emptyset$, each of these p sets must be contracted into an empty set. But since the contraction factors are iid at each stage of the fractal generating process, the probability that any one of these p sets turns out to be empty is equal to the probability that the original set $K_{F,n}$ is empty, and these probabilities are mutually independent. So, we may simplify:

$$q_n = \sum_{p=0}^{N} P(K_{F,n} = \emptyset)^p \cdot P\left(\sum_{i \le N} {R_i^{(n)}}^0 = p\right) = \sum_{p=0}^{N} C_{n,p} \cdot q_n^p.$$

This also holds when $n = 0$.

For each n, define the function $\zeta_n(x) = \sum_{p \le N} C_{n,p} \cdot x^p - x$. For all $n \ge n_0$, $C_{n,p}$ is strictly positive so ζ_n has a root q_n in $[0, 1)$. The same holds for $n = 0$. Since $K_{F,n} \subset K_{F,n+1} \subset K_F$, we have $q_n \ge q_{n+1} \ge q_0$, so $q_\infty \equiv \lim_{n\to\infty} q_n \ge q_0$.

Additionally, for all p, $C_{n,p} \to C_{0,p}$ as $n \to \infty$ since $R_i^{(n)} \to R_i$. This implies that $\zeta_n \to \zeta_0$ uniformly on $[0, q_{n_0}]$. Since $\zeta_n \to \zeta_0$ and $q_n \to q_\infty$, $\zeta_n(q_n) \to \zeta_0(q_\infty)$. But $\zeta_n(q_n) = 0$, so $\zeta_0(q_\infty) = 0$.

We observe that since $C_{0,p} > 0$ at least for $p = 1$, it holds that ζ_0 is either a strictly convex function or it is linear with a negative slope; in either case, it has a unique root in $[0, 1)$. This means $q_\infty = q_0$, so in fact $q_n \to q_0$. This implies that almost surely, $K_F \neq \emptyset$ iff, for some $n \ge n_0$, $K_{F,n} \neq \emptyset$.

Now, for each $n \in \mathbb{N}$, let

$$\phi_n(t) = \mathbb{E}\left[{R_1^{(n)}}^t + \cdots + {R_N^{(n)}}^t\right],$$

and for all $n \ge n_0$, let s_n be the unique solution to $\phi_n(t) = 1$, which exists by Theorem 4.5; this theorem also guarantees that ϕ_n is strictly decreasing. Since $R_i^{(n)} \uparrow R_i$, we also have $\phi_n \le \phi_{n+1} \le \phi$. Thus, if $\phi_n(s_n) = \phi_{n+1}(s_{n+1}) = \phi(s)$, it is clear that $s_n \le s_{n+1} \le s$. This means $s_\infty \equiv \lim_{n\to\infty} s_n \le s$. Conversely,

$$\phi_n(s_n) \ge \phi_n(s_\infty) = \mathbb{E}\left[\sum_{i \le N} {R_i^{(n)}}^{s_\infty}\right] = \mathbb{E}\left[\sum_{i \le N} R_i^{s_\infty} \mathbb{1}_{\{R_i \ge 1/n\}}\right] \uparrow \mathbb{E}\left[\sum_{i \le N} R_i^{s_\infty}\right]$$

by monotone convergence. The first term, $\phi_n(s_n)$, equals 1; and the final term equals $\phi(s_\infty)$. This implies that $\phi(s) = 1 \ge \phi(s_\infty)$, and so $s_\infty \ge s$ since ϕ is decreasing. Therefore, $s_n \to s$.

By Theorem 4.7, the quantity

$$M_{s_n}^{(n)} \equiv \lim_{n\to\infty} \sum_{|\sigma|=n} D_{n,\sigma}^{s_n}$$

has finite moments of all orders. Additionally, since the sequence $\{\sum_{|\sigma|=n} D_{n,\sigma}^{s_n}\}_{n=1}^\infty$ is a martingale (by Theorem 4.6), $\mathbb{E}[M_{s_n}^{(n)}] = \mathbb{E}[D_{n,1}^{s_n} + \cdots + D_{n,N}^{s_n}] = \operatorname{diam}(K)^{s_n}$, which is strictly positive.

For every $n \ge n_0$, the stochastic IFS (F_n, μ_n) satisfies all the conditions of Theorem 4.3 and the additional conditions of Theorem 4.13. So, suppose $t < s$ and $K_F \neq \emptyset$. There must exist some $n \ge n_0$ such that $t < s_n$ and $K_{F,n} \neq \emptyset$. Let γ_n be the random measure for the stochastic IFS (F_n, μ_n), as constructed in Theorem 4.8.

Since $\mathbb{E}[M_{s_n}^n] > 0$, Theorem 4.12 implies that $\gamma_n(K_{F,n}) > 0$. But $K_F \supset K_{F,n}$, and thus $\gamma_n(K_F) > 0$. So, by Theorem 4.13, $\mathscr{H}^t(K_F) = \infty$ for all $t < s$; this implies that $\dim_{\mathscr{H}}(K_F) \ge s$ almost surely whenever $K_F \neq \emptyset$, so we are done. $\square$

5. Measuring Random Fractals: Special Cases

The general result of Theorem 4.3 has some nice special cases that we round out this paper by highlighting.

Example 5.1. The most instructive example of a random fractal is the random equivalent of a Cantor set. We show that almost surely a random Cantor set has Hausdorff dimension $(\sqrt{17}-3)/2$.

The way to generate a random Cantor set would be to, at each step, split every interval into two sub-intervals, choosing the contraction factors uniformly, without overlap. The first few stages of this process are pictured below:

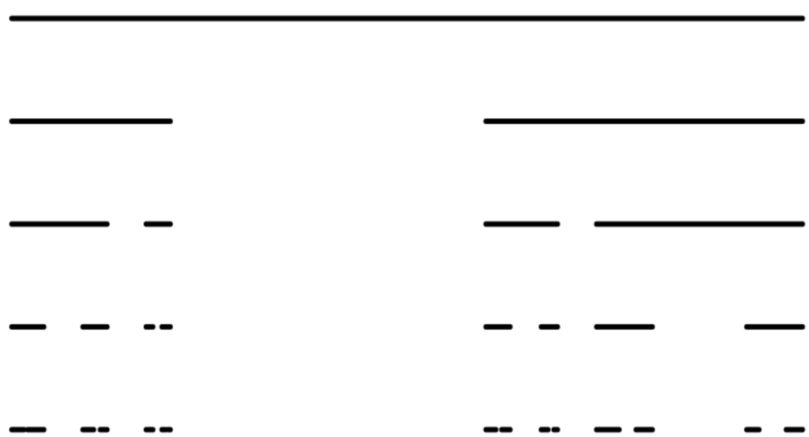

Figure 1: First five stages in the construction of a random Cantor set

To determine the dimension of the resultant Cantor set, we first construct a corresponding stochastic IFS. Let u be the unique uniformly distributed probability measure on the set $S = \{(a,b) : 0 < a \leq b < 1\}$. Now, define the continuous map $g : S \to \tilde{\mathscr{S}}(\mathbb{R})^2$ by

$$g : k \mapsto (F_1(x) = ax, F_2(x) = bx + 1 - b).$$

Then, μ is the pushforward measure of u under the map g. This gives us the stochastic IFS needed to generate a random Cantor set. (It can be checked that there are no concerns with overlapping.)

The Hausdorff dimension of this random set is determined by the solution to the equation $\mathbb{E}[R_1^s + R_2^s] = 1$. We go through the calculations:

$$\begin{aligned}
1 = \mathbb{E}[R_1^s + R_2^s] &= \int a^s + b^s \; du(a,b) \\
&= 2 \cdot \int_0^1 \int_0^{1-a} a^s + b^s \; db \; da \\
&= 2 \cdot \int_0^1 \left[a^s \cdot b + \frac{b^{s+1}}{s+1} \right]_0^{1-a} \; da \\
&= 2 \cdot \int_0^1 a^s - a^{s+1} + \frac{1}{s+1}(1-a)^{s+1} \; da.
\end{aligned}$$

Using a u-substitution, we obtain

$$
\begin{aligned}
&= 2 \cdot \int_0^1 a^s - a^{s+1}\ da + \frac{2}{s+1} \int_0^1 c^{s+1}\ dc \\
&= 2 \cdot \left(\frac{1}{s+1} - \frac{1}{s+2} \right) + \frac{2}{s+1} \cdot \frac{1}{s+2} \\
&= \frac{4}{(s+1)(s+2)}
\end{aligned}
$$

$$
\implies (s+1)(s+2) = s^2 + 3s + 2 = 4 \implies s = \frac{\sqrt{17} - 3}{2}.
$$

Therefore, the Hausdorff dimension of a random Cantor set is $(\sqrt{17} - 3)/2$ with probability 1, so long as the above construction is followed.

The exact same method can be used to randomize the generation of a large class of common fractals, such as the four-corner Cantor set, Sierpinski gasket, von Koch snowflake, etc. As long as the distribution of the contraction factors is known and iid, their Hausdorff dimensions can be determined almost surely.

Example 5.2. Another interesting application of the above results is Mandelbrot percolation. The setup of the model is as follows.

Fix a number $k \geq 2$ and $p \in (0, 1)$. Color the unit square blue and divide it into k^2 squares of side length $1/k$ in the obvious way. For each of the k^2 squares, leave it blue with probability p and color it white with probability $1 - p$, independently of all other squares. Repeat the above operation on each blue square: divide it into k^2 squares, and leave each of the smaller squares blue with probability p and color it white with probability $1 - p$, independently of all the other squares. If we recursively implement the same process, we obtain a random blue set in the limit; we are interested in the dimension of this blue set.

The first six stages in this construction are pictured below, in the $k = 2$ case:

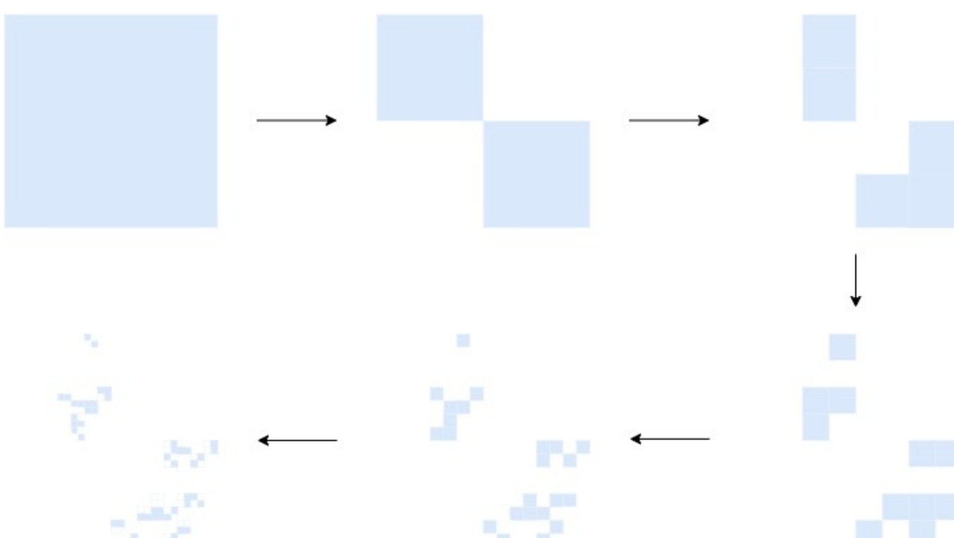

Figure 2: First six stages in the Mandelbrot percolation process

Here we let $K = [0,1]^2$, and define k^2 individual measures $\mu_1, \ldots, \mu_{k^2}$ on $\tilde{\mathscr{S}}(\mathbb{R}^2)$, one for each square sub-division. Since each square is either colored blue or colored white, each function can only take on two possible realizations, both with equal measure; so the measure has two atoms that, together, make up a full measure set. For example, μ_1 corresponds to the first square, so

$$\mu_1(F(x,y) = (\tfrac{1}{k}x, \tfrac{1}{k}y)) = p, \quad \mu_1(F(x,y) = (4,4)) = 1 - p.$$

(Here, $(4,4)$ was chosen as an arbitrary point not in K.) All k^2 measures are defined similarly; and the measure on $\tilde{\mathscr{S}}(\mathbb{R}^2)^N = \tilde{\mathscr{S}}(\mathbb{R}^2)^{k^2}$ is defined as the product measure of $\mu_1, \ldots, \mu_{k^2}$. Now, every r-Lipschitz factor $R_1, \ldots, R_{k^2}$ has an identical distribution, independently of all others:

$$R_i = \begin{cases} \frac{1}{k} & \text{with probability } p, \\ 0 & \text{with probability } 1-p. \end{cases}$$

So, the dimension of the resultant blue set is given by the solution to the equation $\mathbb{E}[R_1^s + \cdots + R_{k^2}^s] = 1$. We solve it out:

$$\begin{aligned} 1 &= \mathbb{E}[R_1^s + \cdots + R_{k^2}^s] \\ &= k^2 \cdot \mathbb{E}[R_1^s] = k^2 \cdot \left(p \cdot \frac{1}{k^s} + (1-p) \cdot 0\right) \\ &= p \cdot k^{2-s} \\ \implies 0 &= \ln p + (2-s) \ln k \\ \implies s &= 2 + \frac{\ln p}{\ln k}. \end{aligned}$$

It is clear that if we were to generalize this $\mathbb{R}^d$ for $d \geq 2$, the dimension would be $s = d + \ln p / \ln k$.

Example 5.3. The framework from Section 4 can also be used to study non-random fractals, simply by turning the measure μ into a δ-measure. Then, the resultant measure P is itself a δ-measure, meaning there is one 'stochastically' μ-self-similar set in its support.

For example, if $K = [0,1]$, $N = 2$, and μ is a point mass such that the double $(F_1(x) = \frac{1}{3}x, F_2(x) = \frac{1}{3}x + \frac{2}{3})$ has full measure, the resultant 'random set' is guaranteed to be the middle-thirds Cantor set; and its Hausdorff dimension can be calculated to be $\ln 2 / \ln 3$.

Clearly, the study of stochastic self-similarity has a large variety of applications and enables the study of a much wider class of fractals than the study of regular self-similarity.

Acknowledgments

I would like to thank my REU mentor, Iqra Altaf, for her invaluable assistance throughout this project. Her mathematical expertise and crystal-clear expository style were of more help to me than I could put in words. Additionally, I thank Professor Peter May of UChicago for organizing the REU and its lectures. In particular, I'd like to shout out the following speakers for delivering wonderful lectures grounded in probability theory and/or analysis: Iqra Altaf, Marianna Csörnyei, Dannin Eccles, Agustin Esteva, Ewain Gwynne, Subhashish Mukherjee, Mahnav Petersen, and Jinwoo Sung.

References

[1] Hutchinson, John. "Fractals and Self-Similarity." Indiana University Mathematics Journal 30, no. 5 (1981): 713. https://doi.org/10.1512/iumj.1981.30.30055.

[2] Graf, Siegfried. "Statistically Self-Similar Fractals." Probability Theory and Related Fields 74, no. 3 (September 1987): 357–92. https://doi.org/10.1007/bf00699096.

[3] Mauldin, Daniel, and S.C. Williams. "Random Recursive Constructions: Asymptotic Geometric and Topological Properties." Transactions of the American Mathematical Society, vol. 295, no. 1 (January 1986): 325–45. https://doi.org/10.1090/s0002-9947-1986-0831202-5.

[4] Mattila, Pertti. Geometry of Sets and Measures in Euclidean Spaces. Cambridge University Press, 1999.