Exploration of the Fractal Dimension of Perlin Noise

Using its Connection to Fractional Brownian Motion

Nicholas Harsell

MATH448 – Numerical Analysis
James Madison University,
Harrisonburg, VA.

Spring 2026

Outline

Fractal Dimension Overview

What is Fractal Dimension?

Fractal dimension is a concept that generalizes the idea of dimensionality to non-integer values.

The fractal dimension is usually used to describe either the "self-similarity" of a shape/image, or the "space-filling ability".

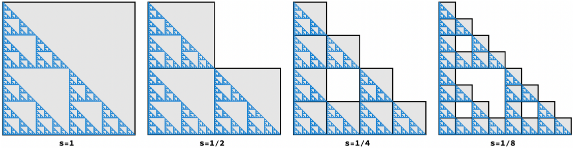

Figure: A visualization of the "box-counting method" with the Sierpinski triangle [2] (next slide).

Box-Counting Method

A standard explanation uses the box-counting method:

1. Place boxes of side length ε across the image, counting each box in which the shape/image is present.
2. Do this as $\varepsilon \to 0$. Theoretically,

$$N(\varepsilon) \sim \varepsilon^{-D},$$

where $N(\varepsilon)$ is the number of boxes counted for side length ε. Then, we find a box-counting dimension D of

$$D = -\log_{\varepsilon}(N) = \frac{\log N(\varepsilon)}{\log(1/\varepsilon)}.$$

Methods of Estimating Dimension (2D)

$2D$ counting methods will find a dimension $2 \leq D \leq 3$ (instead of box-counting, think "cube-counting"). The two methods we use are the following:

1. Power Spectrum (Fourier) method: uses a Fast Fourier Transform and compares the log(power) to log(frequency), finding the line of best fit between the two. Mathematical explanation can be found in section 4.
2. Variogram method: for our usage, Variogram uses Monte Carlo simulation to estimate the difference in pixel intensities of an image based on their distance from each other.

Perlin Noise

Random Noise

- Mathematical noise is a controlled, pseudo-random function designed to simulate natural variation. It provides a way to generate complex, chaotic-looking data that is actually predictable and repeatable based on a given input coordinate.
- Gradient noise generates organic, flowing patterns by placing random vectors on a grid and smoothly interpolating the spaces between them, making it the standard method for simulating natural phenomena like clouds or terrain.

What is Perlin Noise?

- Perlin noise is a gradient noise generation function that generates smooth transitions from high-to-low pixel intensities by interpolating the randomly generated gradient vectors at each pixel. The specific interpolation function and further explanation can be found in [3].
- Most commonly used in procedural texture/terrain generation (Perlin noise is the algorithm used to generate Minecraft terrain!)

The method was initially invented by Ken Perlin [7].

Standard Noise vs. Perlin Noise

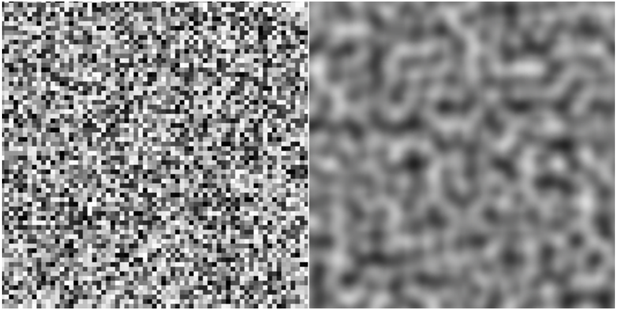

Figure: Standard noise (left) vs. Perlin noise (right)

Initial Experiments

Dimension Across Persistence

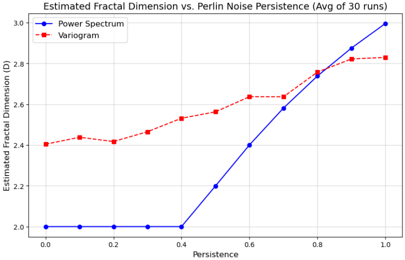

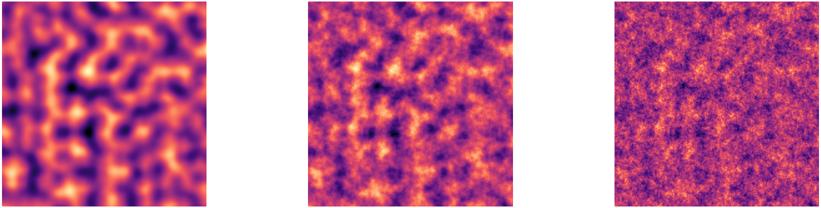

Figure: Dimension estimate w/ persistence values 0.1, 0.5, and 0.9.

Dimension Across Scale

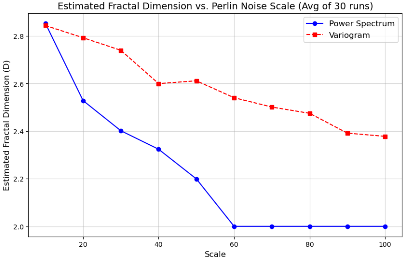

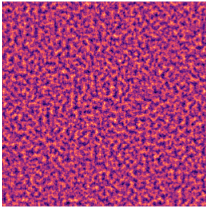

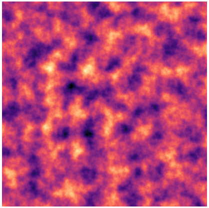

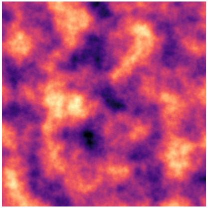

Figure: Dimension estimate w/ scale values 10, 50, and 100.

Dimension Across Octaves

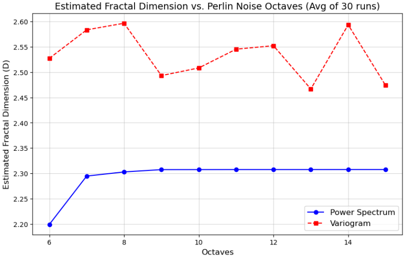

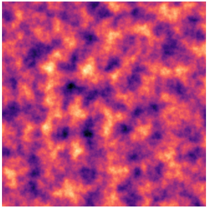

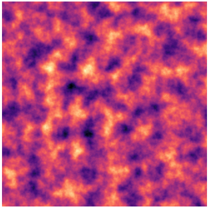

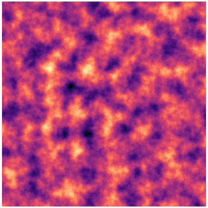

Figure: Dimension estimate w/ octave counts 6, 11, and 16.

Statistical Values (100 Samples)

Standard values for Perlin noise are:

- Persistence: 0.5
- Scale: 50.0
- Octaves: 6

With these values, we find empirical dimensions:

- Power Spectrum Dimension Estimate: 2.195
- Variogram Dimension Estimate: 2.537

Why the difference? Why are the plots of values so inconsistent?

Just Enough Theory

Perlin Noise and fBm

Fractional Brownian motion (fBm) is a generalization of standard Brownian motion in which "generated" values need not be independent of one another.

Layering and rescaling successive octaves of Perlin noise via lacunarity and persistence (covered on slide 19) produces a band-limited approximation of 2D fBm. Consequently, pure fBm probability theory provides an analytical framework to mathematically predict the expected fractal properties of Perlin surfaces.

Hurst Exponent for fBm

The Hurst exponent determines the persistence of fBm through the covariance function

$$\mathbb{E}[B_t^H B_s^H] = \frac{1}{2}\left(t^{2H} + s^{2H} - |t-s|^{2H}\right)$$

For the Gaussian process $\{B_t^H, t \geq s \geq 0\}$, where H is the Hurst exponent [4] [5]. For $H = 0.5$:

$$\mathbb{E}[B_t^{0.5} B_s^{0.5}] = \frac{1}{2}\left(t + s - |t-s|\right) = \frac{1}{2}(t + s - t + s) = s.$$

In English: to have Perlin noise approximate standard brownian motion (uncorrelated values), we must find the values that give $H = 0.5$.

Persistence and the Hurst Exponent

- Persistence p: Affects the roughness of Perlin noise/fBm through how much each successive octave contributes to final values.
- Lacunarity ℓ: The scaling factor of each successive octave/layer of Perlin noise.

The standard relationship between p, ℓ, and H is $p = \ell^{-H}$ (see slide 18). Standard lacunarity is $\ell = 2$; thus, to find $H = 0.5$,

$$\begin{aligned} p &= 2^{-0.5} \\ &= \frac{1}{\sqrt{2}} \\ p &\approx 0.7071068. \end{aligned}$$

Statistical Values pt. 2

We now use values of

- Persistence: ≈ 0.7071068
- Scale: 50.0
- Octaves: 6

With these values, we find empirical dimensions:

- Power Spectrum Dimension Estimate: 2.5888
- Variogram Dimension Estimate: 2.6611

There is no significant difference between the two methods now; but why are we off of the expected 2.5?

Expected Outcomes

Frequency f is multiplied by ℓ at each successive octave; for $\ell = 2$, frequency is doubled, and amplitude

$$A(f) \propto f^{-H}$$

is multiplied by p. Therefore, we find for $\ell = 2$,

$$(2f)^{-H} = pf^{-H} \implies \underbrace{p = 2^{-H} \quad \text{or} \quad H = -\log_2(p)}_{\text{(Relationship from before!)}}.$$

From this, we find the theoretical dimension D_T for given persistence p to be

$$D_T(p) = 3 + \log_2(p),$$

$$\implies D_T(0.5) = 2, \qquad D_T(\sqrt{2}/2) = 2.5.$$

Method 1: Power Spectrum (Fourier)

The Power Spectrum method works as follows:

1. Compute FFT, $F(u, v)$
   - u and v represent spatial frequencies in the x and y directions
2. Calculate Power spectrum as $P(u, v) = |F(u, v)|^2$
3. Assuming our image is isotropic, we evaluate how P changes as frequency gets higher:
   - $r = \sqrt{(u - u_0)^2 + (v - v_0)^2} : \quad (u_0, v_0) =$ center
   - $P(r) = \frac{1}{N_r} \sum_{u,v \in \text{ring}(r)} P(u, v)$
4. $P(r) \propto r^{-\beta}$

Which we then use for $H = \frac{\beta - 2}{2}$ [1] and $D = \frac{8 - \beta}{2}$ [5].

The Satisfying Answer (Power Spectrum)

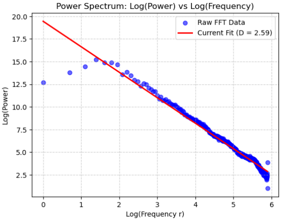

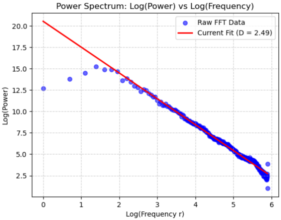

log(r) when $r \approx 0$ yields a biased slope, $-\beta$. We see this in the plot on the left:

$$2.59 = \frac{8-\beta}{2} \implies \beta = 2.82.$$

When we ignore the first 5-10 biased terms, our estimate then gives $\beta \approx 3$, giving a near-perfect dimension estimate of 2.49!

Method 2: Variogram Estimate

Standard variogram for a 512×512 image would require evaluating $\approx 3.4 \times 10^{10}$ pairs of pixels, so we estimate through Monte Carlo. For each n in `number_of_samples`, choosing a pair of two points $p_1 = (x_1, y_1)$ and $p_2 = (x_2, y_2)$:

1. $d(p_1, p_2) = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}$
2. We estimate the theoretical Variogram through

$$2\gamma(d) \approx \frac{1}{n} \sum_{k=0}^{n-1} (Z(p_{1,k}) - Z(p_{2,k}))^2 = V(d)$$

3. $V(d) \propto d^{2H}$ [6].

The Unsatisfying Answer

We are off of the expected $D_T = 2.5$ due to estimation error and the finite nature of our image (recall that Variogram uses Monte Carlo Simulation). To find a more accurate estimate, D_E, we must get more data.

Recalculating our Variogram with a 4096×4096 image ($16\times$ larger) and $50,000$ trials ($10\times$ more), we get:

- $D_E = 2.0$ with $p = 0.5$ ($D_T = 2.0$)
- $D_E = 2.51$ with $p \approx \frac{1}{\sqrt{2}}$ ($D_T = 2.5$)

(Plots on Next Slide)

Variogram Diagnosis Images

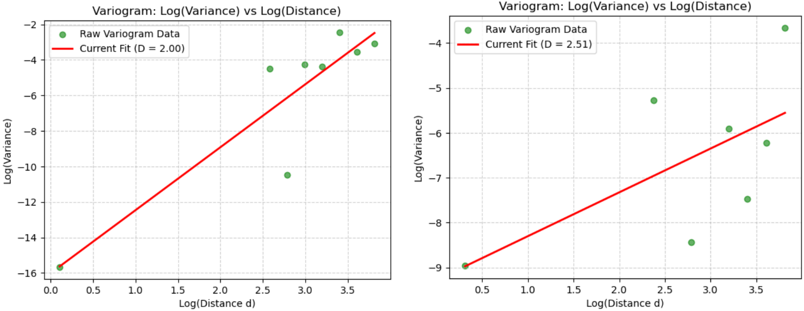

The small amount of values is due to the bin size required for accurate estimates. More trials in the Monte Carlo Simulation can allow more bins.

Conclusion

Theory Takeaways

- Using the relationship $p = \ell^{-H}$, we can use different values of p and ℓ to generate Perlin noise of any theoretical dimension $2 \leq D_T \leq 3$.
- We can approximate standard Brownian motion through $\ell = 2$ and $p = \frac{1}{\sqrt{2}}$.
- Can use this theory to generate pseudo-random terrain of any specific roughness and scale.

Numerical Method Takeaways

- Power Spectrum tends to be a significantly faster and more numerically stable estimate of fractal dimension compared to the Variogram method, but can take longer to converge to true values if extreme accuracy is necessary.
- Both methods can be used to estimate the fractal dimension of 2D noise, though each require their own considerations and numerical analysis.

Thank you!

Source code available at:
nichhname/fractal-dimension

Michael F. Barnsley, Robert L. Devaney, Benoit B. Mandelbrot, Heinz-Otto Peitgen, Dietmar Saupe, and Richard F. Voss.
The Science of Fractal Images.
Springer New York, 1988.

Harlan Brothers.
A Gentle Introduction to Box-Counting — sciencespectrumu.com.
https://sciencespectrumu.com/a-gentle-introduction-to-box-counting-202e0106c375, 2024.

Flafla2.
Understanding Perlin Noise — adrianb.io.
https://adrianb.io/2014/08/09/perlinnoise.html, 2014.

Andrey Nikolaevich Kolmogorov.
Wienersche spiralen und einige andere interessante kurven im hilbertschen raum.
C. R. (Doklady) Acad. Sci. URSS (NS), 26(2):115–118, 1940.

Benoit B. Mandelbrot and John W. Van Ness.
Fractional brownian motions, fractional noises and applications.
SIAM Review, 10(4):422–437, 1968.

Alex P. Pentland.
Fractal-based description of natural scenes.
IEEE Transactions on Pattern Analysis and Machine Intelligence, 6(6):661–674, 1984.

Ken Perlin.
An image synthesizer.
ACM SIGGRAPH Computer Graphics, 19(3):287–296, 1985.