Fractals: Self-Similarity and Fractal Dimension

Math 198, Spring 2013

Background

Fractal geometry is one of the most important developments in mathematics in the second half of the 20th century. Fractals are central to understanding a wide variety of chaotic and nonlinear systems, and so have many applications in the sciences. However, they are also beautiful objects in their own right; below are pictures of two classic fractals, the *Mandelbrot set* and the *Julia set* (produced using the program *Ultra Fractal*).

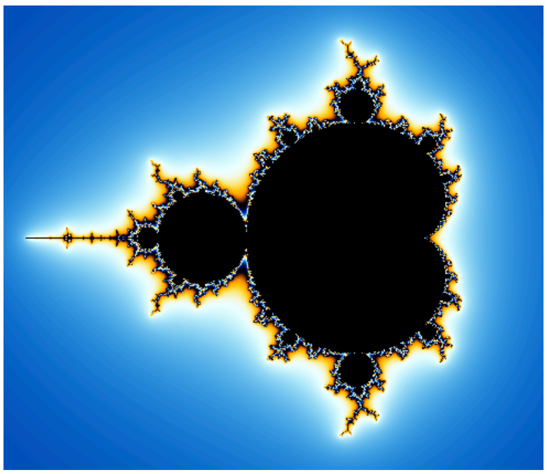

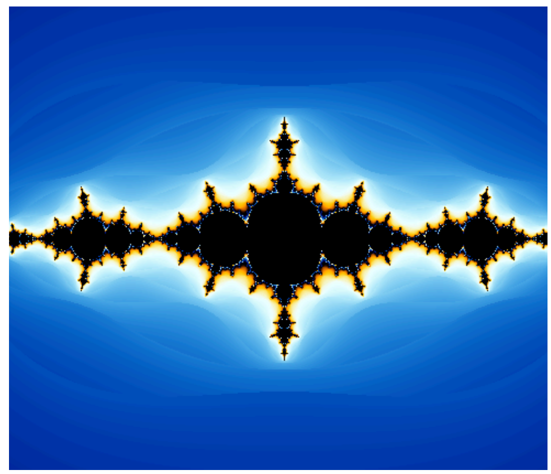

What is a fractal? A fractal is a geometric object whose *fractal* dimension is larger than its *topological* dimension. This is not a very well-defined notion – there are many ways of measuring dimension (both fractal and topological), and they do not all agree. But, intuitively, a fractal is a geometric object which is "infinitely complicated" – while it may, topologically, be a curve or a surface, no matter how much it is magnified it will never "smooth out" to resemble an Euclidean space. The various notions of fractal dimension attempt to quantify this complexity.

Many fractals also have a property of *self-similarity* – within the fractal lies another copy of the same fractal, smaller but complete. In this project we will study the simplest, and best known, fractals with this property, the *strictly self-similar* fractals. We will show how to describe and create these fractals, and how to measure their fractal dimension using the *similarity dimension*.

Strictly Self-Similar Fractals

A geometric figure is **self-similar** if there is a point where every neighborhood of the point contains a copy of the entire figure. For example, imagine the figure formed by inscribing a square within another square, rotated by 45°. Then inside the inner square, inscribe another square in the same manner, and so on *ad infinitum*. Of course, we can't really draw this figure, since it contains infinitely many nested squares, but an approximation of the result is shown below:

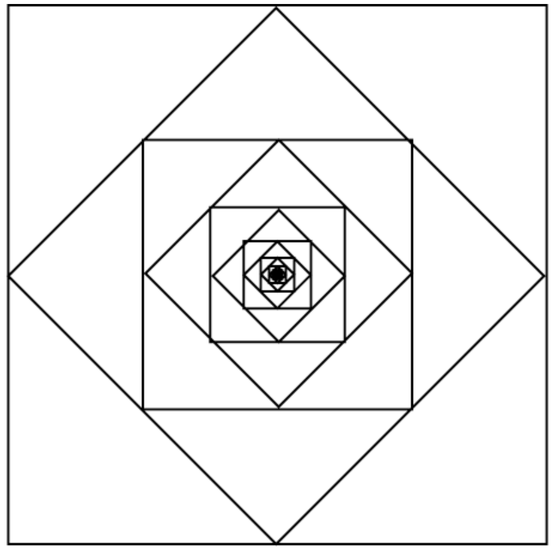

This figure is self-similar at the center of the squares – every ball around the center, no matter how small, will contain a complete copy of the entire figure. However, this is the only point where the figure is self-similar – at most points, the figure looks, at some magnification, like either a straight line or a corner where four lines meet, so this is not complicated enough to be a fractal.

A figure is **strictly self-similar** if it is self-similar at *every* point. Equivalently, this means that the figure can be decomposed into some number of disjoint pieces, each of which is an exact copy of the entire figure. What does such a figure look like? Let's look at several classical examples. As with the inscribed squares above, these examples are really the *limits* of some repeated, or *iterative*, process, which we imagine repeating infinitely often.

Cantor Set

The Cantor set was first described by German mathematician Georg Cantor in 1883, and is the foundation for many important fractals. The description of the set is deceptively easy. Begin with the unit interval [0,1], which is represented below as a line segment (stage 0). Delete the middle third, which is the open interval (1/3, 2/3), leaving the two closed subintervals [0, 1/3] and [2/3, 1]. Now repeat this process on each of these subintervals, leaving 4 subintervals, and continue repeating this process on each new set of smaller subintervals forever. The figure below shows the first few stages of this construction:

The Cantor set is the set of points remaining when all of these deletions of subintervals have taken place. This is more than you might expect – for example, all the endpoints of all the closed intervals at every stage of the process are in the Cantor set. In fact, there are just as many points in the Cantor set as there were in the original unit interval![1] Despite this, any point of the interval [0, 1] that is *not* in the Cantor set is the center of an open interval which contains *no* points of the Cantor set – we say that the Cantor set is *nowhere dense*. So the Cantor set is simultaneously as large as the entire unit interval, and so small that any point not in the set can be separated from it with an open interval. It is truly an amazing mathematical object!

However, we are most interested in the fact that the Cantor set is strictly self-similar: the points in the Cantor set are all either in [0, 1/3] or [2/3, 1], and the subset of the Cantor set in either of these intervals is an exact copy of the entire set. So the Cantor set can be decomposed into two disjoint pieces, each of which is itself a Cantor set. In fact, at any stage n, the Cantor set can be decomposed into 2^n disjoint pieces, each an exact copy of the entire set.

Koch Curve

Swedish mathematician Helge von Koch introduced the *Koch curve* in 1904, as an example of a curve that is continuous but nowhere differentiable. It is the result of a simple geometric construction that is easily generalized to give a wide variety of examples. Begin with a line segment in the plane – for example, the interval [0, 1] along the x-axis (stage 0). On the middle third of this interval, construct an equilateral triangle upward (so each side of the triangle has length 1/3, and one side is the interval [1/3, 2/3]), and remove the base (the open interval (1/3, 2/3)). The result is four line segments of length 1/3, arranged as shown below (stage 1). For stage 2, perform the same construction on each of these four line segments so that the new segments protrude outward from those of stage 1 (see below) and continue to repeat the process indefinitely on all the smaller line segments formed at each stage. A few stages of this process are shown below:

[1] To see this, observe that any point of the Cantor set can be described as an infinite sequence of L's and R's, indicating whether it was in the Left or Right subinterval of the interval containing it in the previous stage. For example, 20/27 corresponds to RLRLLLLL... (L's forever). Any such sequence can be converted to a "binary decimal" of 0's and 1's by replacing L by 0 and R by 1, so 20/27 would correspond to the decimal 0.1010000... This maps the points of the Cantor set onto the set of binary decimals. But since every number between 0 and 1 has a representation as a binary decimal (writing its decimal expansion in base 2), this means there are at least as many points in the Cantor set as in the entire unit interval [0, 1]!

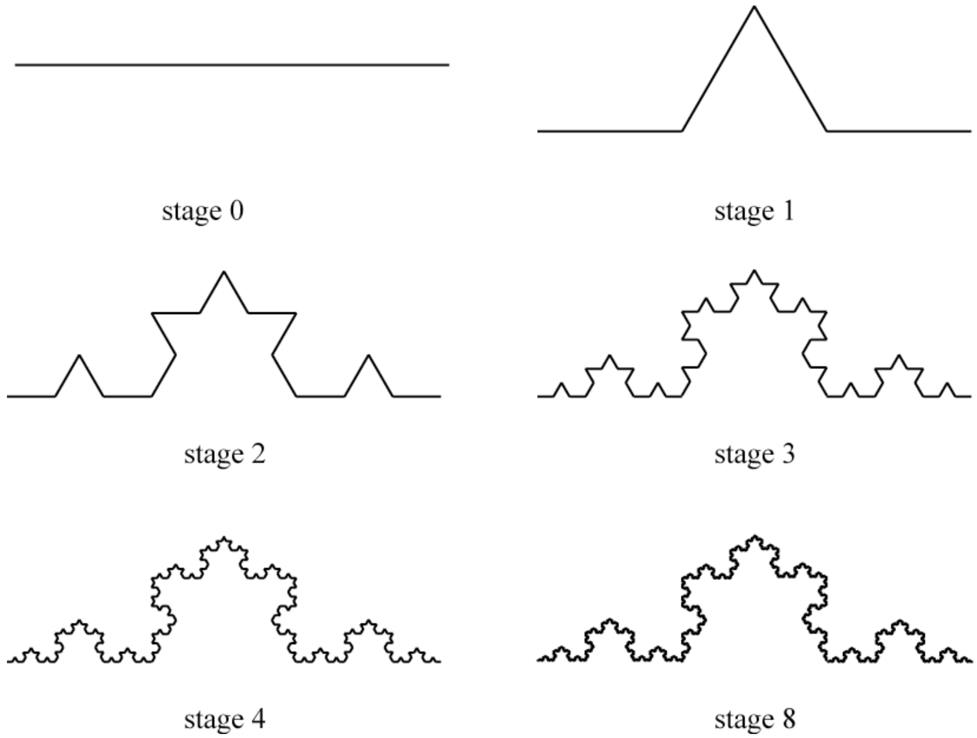

<u>Exercise 1:</u> Explain why the Koch curve is strictly self-similar.

Sierpinski Triangle

Our next example is the Sierpinski triangle, introduced in 1916 by the Polish mathematician Waclaw Sierpinski. Begin with a solid equilateral triangle. Joining the midpoints of the three sides gives another equilateral triangle – remove it. You will be left with three equilateral triangles, with each pair meeting at a corner. Repeat the process with each of these equilateral triangles, and continue to repeat the process on each new smaller solid triangle left at each stage. Several early stages of the construction are shown below.

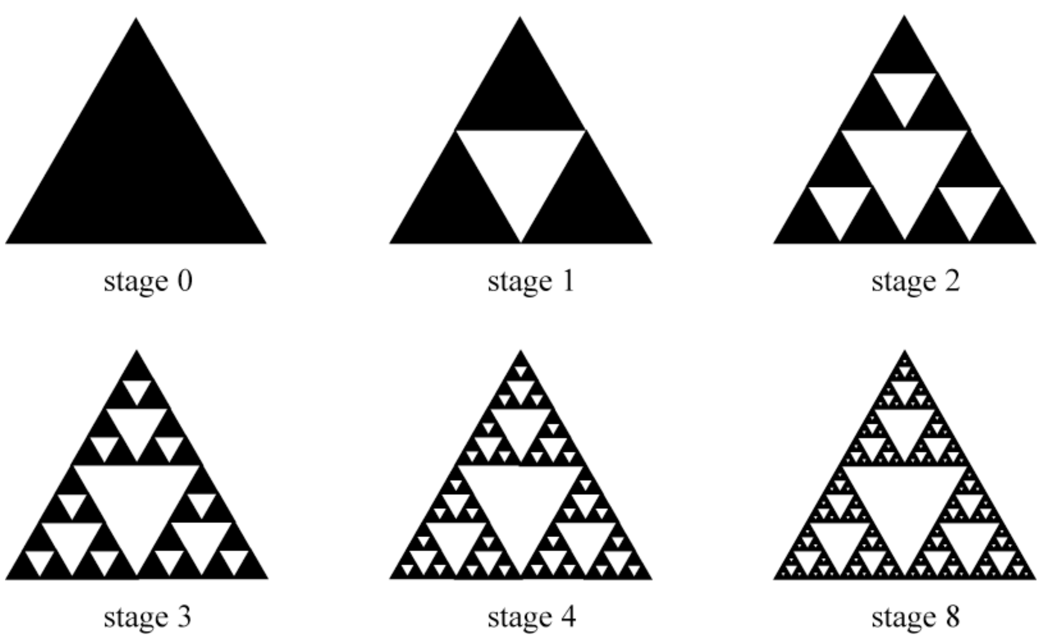

<u>Exercise 2:</u> Explain why the Sierpinski triangle is strictly self-similar.

Creating Fractals with Similarity Transformations

A **similarity transformation** is a combination of *scaling, rotations, translations* and *reflections*. The result of performing a similarity transformation on a figure is an exact copy of the figure, only resized (and possibly rotated or reflected). So all angles are preserved, and the shape of the figure is the same.

We can construct any strictly self-similar fractal by transforming a geometric figure using a combination of similarity transformations. For example, the Cantor set involves performing two similarity transformations on the interval [0, 1]: scaling it by a factor of 1/3 to produce the interval [0, 1/3], and scaling it by a factor of 1/3 combined with a translation of 2/3 units to the right to produce the interval [2/3, 1]. We call these two transformations t_1 and t_2. Given any figure A, we can create a new figure $T(A) = t_1(A) \cup t_2(A)$. For the Cantor set we begin with A_0 = [0, 1], and let $A_1 = T(A_0)$. Then, in general, $A_{n+1} = T(A_n)$, and the Cantor set itself is A_∞. This means that the Cantor set can also be defined as the *fixed set* of the transformation T.

<u>Exercise 3:</u> The Koch curve and Sierpinski triangle can each be described as the fixed set of a transformation made up of similarity transformations. In each case, find the transformation.

Using similarity transformations gives us a new way to create many new strictly self-similar fractal designs. Start with any basic figure, and choose a scaling factor less than 1. Make several smaller copies of the figure (all using the same scaling factor in every direction), and arrange them in some pattern. For the moment, make sure that the copies don't overlap (though they can touch at the corners or sides), or the resulting fractal won't be strictly self-similar; also the new arrangement should be roughly the same size as the original figure. Then repeat the process. This can be carried out quite easily with a computer graphics program, or even with a photocopier and scissors.

Just as we have previously considered a symmetry of a set of points in the plane to be a *rigid motion* that leaves the set looking the same, we can now consider a symmetry of a strictly self-similar fractal to be any *similarity transformation* that leaves the set looking the same; so our symmetries now include *scaling* as well as the familiar rigid motions.

Sierpinski Carpet

As another example of this process, we will look at another fractal due to Sierpinski. Begin with a solid square. Make 8 copies of the square, each scaled by a factor of 1/3 (both vertically and horizontally), and arrange them to form a new square the same size as the original, with a hole in the middle. Repeat the process on each of the smaller scaled squares, and carry out the iterated process forever. The first three stages are shown below:

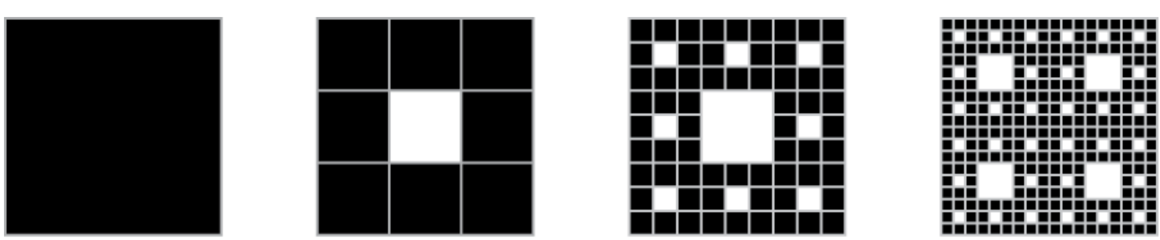

stage 0 stage 1 stage 2 stage 3

Menger Sponge

As a final classical example, we will look at a 3-dimensional analogue of the Sierpinski carpet. Begin with a solid cube, and make 20 smaller copies by scaling the cube by a factor of 1/3. Arrange the smaller cubes to form the "skeleton" of a cube the size of the original (leaving out a cube in the middle, and a cube in the middle of each side). Repeat the process on each of the smaller cubes, and continue the process indefinitely. The first three stages are shown below (image courtesy of [So]):

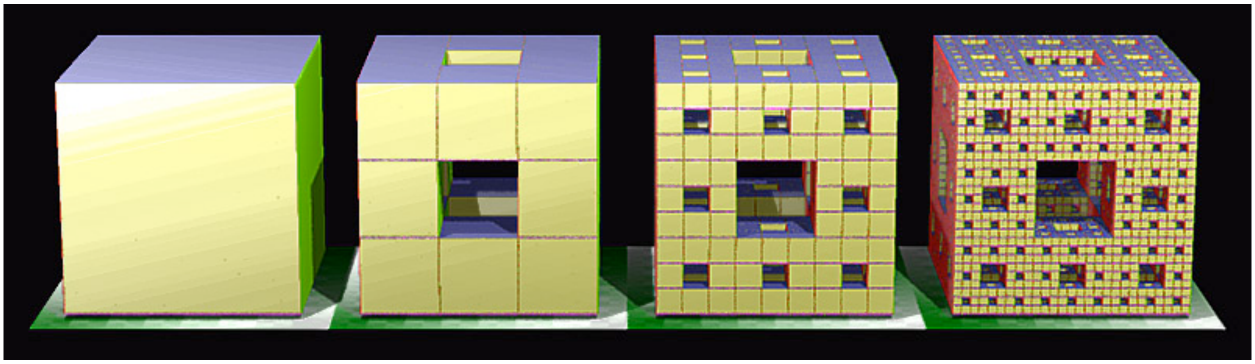

Exercise 4: Create your own strictly self-similar fractal design by following the steps below.

(a) Begin with any figure (e.g., a drawing, photograph, signature), and plan an arrangement of copies of the figure as was done in the examples above. Using a computer or photocopier, perform several iterations of the reducing/arranging process to get a good approximation of the fractal (generally, by 10 iterations or fewer you'll have reached the limit of available resolution).

(b) Repeat part (a) using a different original design, but the same reduction/arrangement process. How does the resulting fractal compare to your first one?

You should find that the original design is irrelevant – it's the transformation that determines the final fractal.

Fractal Dimension

What distinguishes fractals from classical geometric figures is their *dimension*. The notion of dimension is very familiar, but surprisingly subtle. Intuitively, we know that a line or curve is one-dimensional, a plane or surface is two-dimensional and space is 3-dimensional. Mathematicians consider a figure to be one-dimensional if it can be cut into pieces which each look like a piece of a line, two-dimensional if it can be cut into pieces which look like a piece of a plane, and so forth. But this rough notion of dimension doesn't work for fractals.

Consider the Koch curve. At each stage of the construction, the figure consists of a finite number of one-dimensional line segments, and so is one-dimensional. But the final curve *cannot* be cut into line segments – even if we cut it into infinitely many pieces, if a piece is larger than a single point, then it contains a complete copy of the entire curve. So no matter how small the pieces are, they will not look like pieces of a line, indicating that the Koch curve is *more* than one-dimensional – but neither do they look like pieces of a plane, so it's *less* than two-dimensional. This leads us to look for a new definition of dimension that can have non-integer values, called a *fractal dimension.*

There are several ways to define a fractal dimension (and they are not always equivalent). We will use a definition that is useful for strictly self-similar fractals, called the *similarity dimension*.

This definition is based on the fact that these fractals can be cut into pieces identical to the original fractal, but scaled by some factor. The question is then how to relate the number of pieces the original was figure cut into and the scaling factor. To understand this relationship, we will first consider our simplest geometric objects in dimensions 1, 2, and 3 – the line segment, square, and cube.

Scale each of these figures by a factor of 1/2. How many of the smaller copies do you need to assemble to get the original figure?

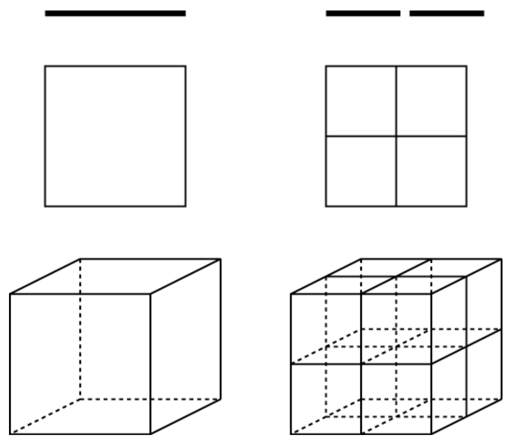

For the line segment, it requires 2 copies, for the square 4, and for the cube 8.

Exercise 5: What if we scale by a factor of 1/3 instead? 1/4? 1/5? How many copies do you need to assemble to obtain each of the original three figures? What do you notice?

You should notice that if you scale by a factor of $1/k$, the line segment requires k copies, the square requires k^2 copies, and the cube requires k^3 copies. This indicates that if you take a d-dimensional figure and scale it by a factor of $1/k$, you will need k^d copies to recreate the original object.

So if we have a strictly self-similar fractal which we can decompose into n pieces, each of which is a copy of the original fractal scaled by a factor of $1/k$, then:

$$n = k^d$$
$$\log(n) = \log(k^d)$$
$$\log(n) = d\log(k)$$
$$d = \frac{\log(n)}{\log(k)}$$

So we can use logarithms to solve for the dimension d. This number d is called the **similarity dimension** of the figure.

Exercise 6: Compute the similarity dimension of (a) the Cantor set (b) the Koch curve (c) the Sierpinski triangle (d) the Sierpinski carpet (e) the Menger sponge and (f) the fractal you created for Exercise 4.

We can also produce fractals with a predetermined similarity dimension. For example, to produce a fractal with dimension 3/2, we need $\log(n)/\log(k) = 3/2$; so we need to have some x

where $n = x^3$ and $k = x^2$.[2] We can choose any value of x we want; if we choose $x = 2$, then $n = 8$ and $k = 4$, so we want to make an arrangement of 8 copies, each scaled by a factor of 1/4. An example of such an arrangement, and the first few stages of the resulting fractal, is shown below; the fractal has dimension 3/2, and so this is called the 3/2-curve:

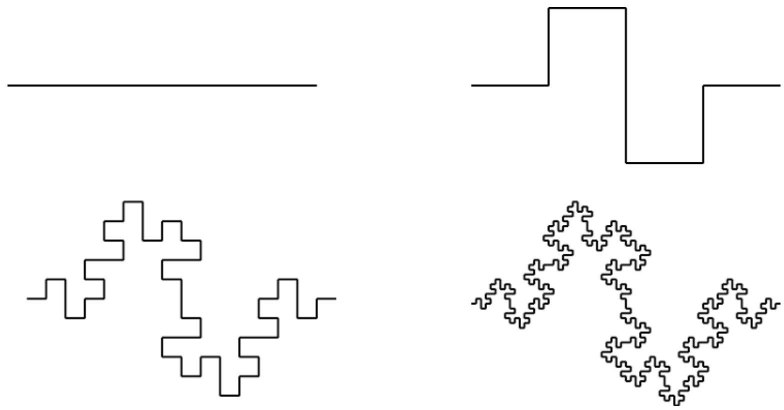

<u>Exercise 7:</u> Design fractals with similarity dimension (a) 4/3 (b) 5/3 and (c) 5/4.

References

[Ma] Mandelbrot, B.: *The Fractal Geometry of Nature*, W.H. Freeman and Co., San Francisco, 1977

[PJS] Peitgen, H.-O., Jürgens, H., and Saupe, D.: *Fractals for the Classroom*, Springer-Verlag, New York, 1992

[So] Solkoll: *Image:Menger sponge (Level 1-4).jpg,* Wikipedia, *http://en.wikipedia.org/wiki/Image:Menger_sponge_%28Level_1-4%29.jpg*. Public domain image

---

[2] If the math isn't obvious here, note that equation $\log(n)/\log(k) = 3/2$ implies that $\log(n) = 3t$ and $\log(k) = 2t$ for some scalar t. So $n = 10^{3t} = (10^t)^3$ and $k = 10^{2t} = (10^t)^2$. If we let $x = 10^t$, then $n = x^3$ and $k = x^2$.