10.2 Logistic Map: Numerics

In a fascinating and influential review article, Robert May (1976) emphasized that even simple nonlinear maps could have very complicated dynamics. The article ends memorably with "an evangelical plea for the introduction of these difference equations into elementary mathematics courses, so that students' intuition may be enriched by seeing the wild things that simple nonlinear equations can do."

May illustrated his point with the ***logistic map***

$$x_{n+1} = rx_n(1 - x_n), \qquad (1)$$

a discrete-time analog of the logistic equation for population growth (Section 2.3). Here $x_n \geq 0$ is a dimensionless measure of the population in the nth generation and $r \geq 0$ is the intrinsic growth rate. As shown in Figure 10.2.1, the graph of (1) is a parabola with a maximum value of $r/4$ at $x = \frac{1}{2}$. We restrict the control parameter r to the range $0 \leq r \leq 4$ so that (1) maps the interval $0 \leq x \leq 1$ into itself. (The behavior is much less interesting for other values of x and r—see Exercise 10.2.1.)

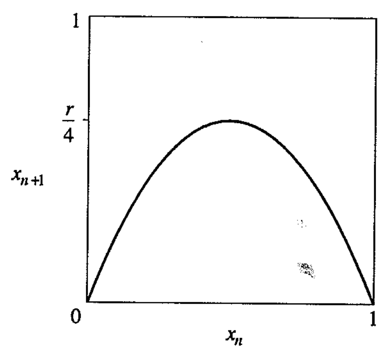

Figure 10.2.1

Period-Doubling

Suppose we fix r, choose some initial population x_0, and then use (1) to generate the subsequent x_n. What happens?

For small growth rate $r < 1$, the population always goes extinct: $x_n \to 0$ as $n \to \infty$. This gloomy result can be proven by cobwebbing (Exercise 10.2.2).

For $1 < r < 3$ the population grows and eventually reaches a nonzero steady state (Figure 10.2.2). The results are plotted here as a ***time series*** of x_n vs. n. To make the sequence clearer, we have connected the discrete points (n, x_n) by line segments, but remember that only the corners of the jagged curves are meaningful.

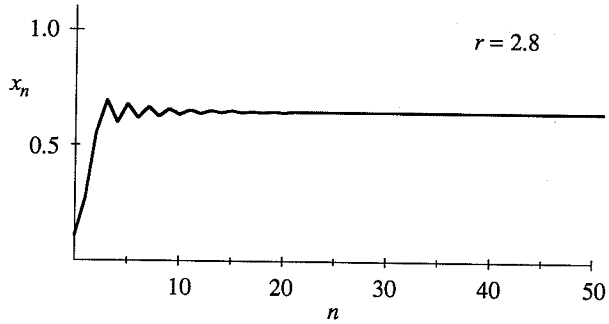

Figure 10.2.2

For larger r, say $r = 3.3$, the population builds up again but now *oscillates* about the former steady state, alternating between a large population in one generation and a smaller population in the next (Figure 10.2.3). This type of oscillation, in which x_n repeats every *two* iterations, is called a ***period-2 cycle***.

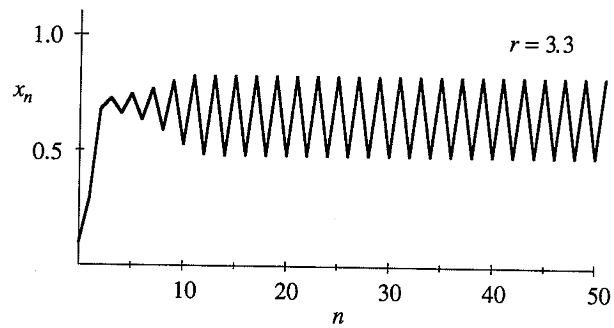

Figure 10.2.3

At still larger r, say $r = 3.5$, the population approaches a cycle that now repeats every *four* generations; the previous cycle has doubled its period to ***period-4*** (Figure 10.2.4).

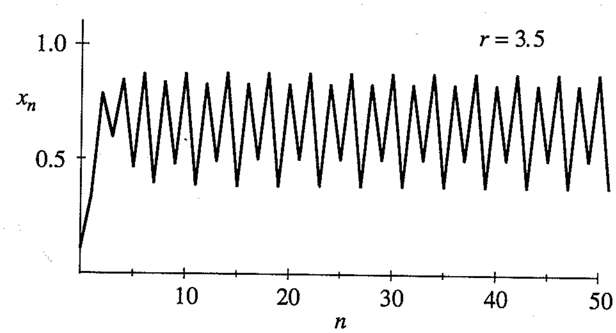

Figure 10.2.4

Further ***period-doublings*** to cycles of period 8, 16, 32, . . . , occur as r increases. Specifically, let r_n denote the value of r where a 2^n-cycle first appears. Then computer experiments reveal that

$r_1 = 3$	(period 2 is born)
$r_2 = 3.449...$	4
$r_3 = 3.54409...$	8
$r_4 = 3.5644...$	16
$r_5 = 3.568759...$	32
$\vdots$	$\vdots$
$r_\infty = 3.569946...$	∞

Note that the successive bifurcations come faster and faster. Ultimately the r_n converge to a limiting value r_∞. The convergence is essentially geometric: in the limit of large n, the distance between successive transitions shrinks by a constant factor

$$\delta = \lim_{n\to\infty} \frac{r_n - r_{n-1}}{r_{n+1} - r_n} = 4.669\ldots.$$

We'll have a lot more to say about this number in Section 10.6.

Chaos and Periodic Windows

According to Gleick (1987, p. 69), May wrote the logistic map on a corridor blackboard as a problem for his graduate students and asked, "*What the Christ happens for* $r > r_\infty$?" The answer turns out to be complicated: For many values of r, the sequence $\{x_n\}$ never settles down to a fixed point or a periodic orbit—instead the long-term behavior is aperiodic, as in Figure 10.2.5. This is a discrete-time version of the chaos we encountered earlier in our study of the Lorenz equations (Chapter 9).

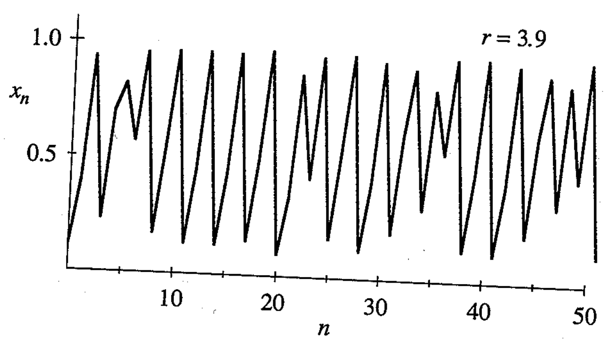

Figure 10.2.5

The corresponding cobweb diagram is impressively complex (Figure 10.2.6).

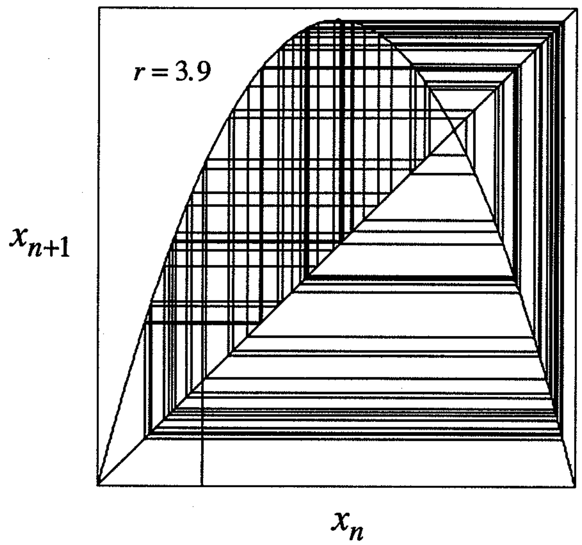

Figure 10.2.6

You might guess that the system would become more and more chaotic as r increases, but in fact the dynamics are more subtle than that. To see the long-term behavior for *all* values of r at once, we plot the ***orbit diagram***, a magnificent picture that has become an icon of nonlinear dynamics (Figure 10.2.7). Figure 10.2.7 plots the system's attractor as a function of r. To generate the orbit diagram for yourself, you'll need to write a computer program with two "loops." First, choose a value of r. Then generate an orbit starting from some random initial condition x_0. Iterate for 300 cycles or so, to allow the system to settle down to its eventual behavior. Once the transients have decayed, plot many points, say $x_{301}, \ldots, x_{600}$ above that r. Then move to an adjacent value of r and repeat, eventually sweeping across the whole picture.

Figure 10.2.7 shows the most interesting part of the diagram, in the region $3.4 \le r \le 4$. At $r = 3.4$, the attractor is a period-2 cycle, as indicated by the two branches. As r increases, both branches split simultaneously, yielding a period-4 cycle. This splitting is the period-doubling bifurcation mentioned earlier. A cascade of further period-doublings occurs as r increases, yielding period-8, period-16, and so on, until at $r = r_\infty \approx 3.57$, the map becomes chaotic and the attractor changes from a finite to an infinite set of points.

For $r > r_\infty$ the orbit diagram reveals an unexpected mixture of order and chaos, with ***periodic windows*** interspersed between chaotic clouds of dots. The large window beginning near $r \approx 3.83$ contains a stable period-3 cycle. A blow-up of part of the period-3 window is shown in the lower panel of Figure 10.2.7. Fantastically, a copy of the orbit diagram reappears in miniature!

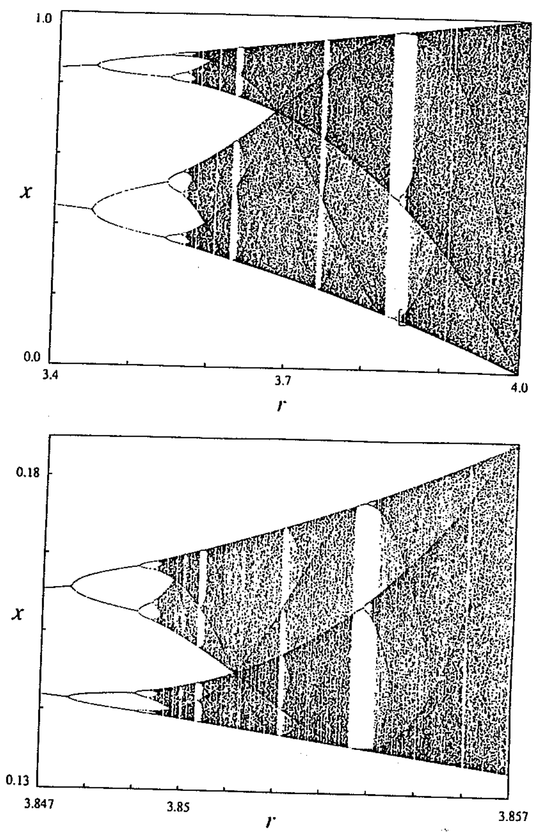

Figure 10.2.7 Campbell (1979), p. 35, courtesy of Roger Eckhardt

10.3 Logistic Map: Analysis

The numerical results of the last section raise many tantalizing questions. Let's try to answer a few of the more straightforward ones.

EXAMPLE 10.3.1:

Consider the logistic map $x_{n+1} = rx_n(1 - x_n)$ for $0 \le x_n \le 1$ and $0 \le r \le 4$. Find all the fixed points and determine their stability.

Solution: The fixed points satisfy $x^* = f(x^*) = rx^*(1 - x^*)$. Hence $x^* = 0$ or $1 = r(1 - x^*)$, i.e., $x^* = 1 - \frac{1}{r}$. The origin is a fixed point for all r, whereas $x^* = 1 - \frac{1}{r}$ is in the range of allowable x only if $r \ge 1$.

Stability depends on the multiplier $f'(x^*) = r - 2rx^*$. Since $f'(0) = r$, the origin is stable for $r < 1$ and unstable for $r > 1$. At the other fixed point, of x_1. Similarly, its second digit is anything other than the second digit of x_2. In general, the nth digit of r is $\bar{x}_{nn}$, defined as any digit other than x_{nn}. Then we claim that the number $r = \bar{x}_{11}\bar{x}_{22}\bar{x}_{33}\cdots$ is not on the list. Why not? It can't be equal to x_1, because it differs from x_1 in the first decimal place. Similarly, r differs from x_2 in the second decimal place, from x_3 in the third decimal place, and so on. Hence r is not on the list, and thus X is uncountable. ■

This argument (devised by Cantor) is called the ***diagonal argument***, because r is constructed by changing the diagonal entries x_{nn} in the matrix of digits $[x_{ij}]$.

11.2 Cantor Set

Now we turn to another of Cantor's creations, a fractal known as the Cantor set. It is simple and therefore pedagogically useful, but it is also much more than that—as we'll see in Chapter 12, the Cantor set is intimately related to the geometry of strange attractors.

Figure 11.2.1 shows how to construct the Cantor set.

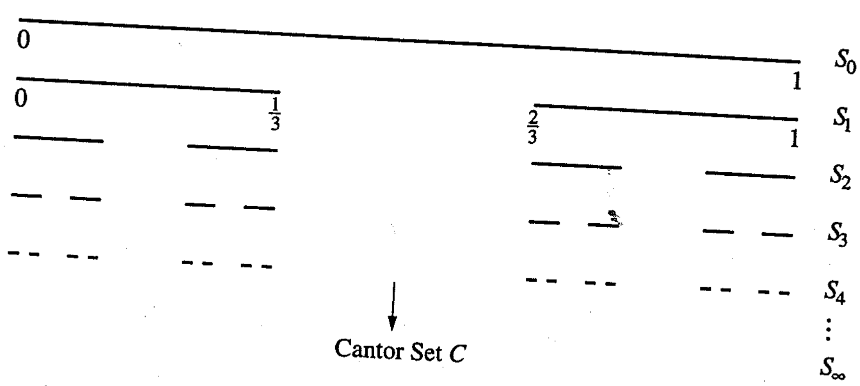

Figure 11.2.1

We start with the closed interval $S_0 = [0,1]$ and remove its open middle third, i.e., we delete the interval $(\frac{1}{3}, \frac{2}{3})$ and leave the endpoints behind. This produces the pair of closed intervals shown as S_1. Then we remove the open middle thirds of *those* two intervals to produce S_2, and so on. The limiting set $C = S_\infty$ is the ***Cantor set***. It is difficult to visualize, but Figure 11.2.1 suggests that it consists of an infinite number of infinitesimal pieces, separated by gaps of various sizes.

Fractal Properties of the Cantor Set

The Cantor set C has several properties that are typical of fractals more generally:

1. *C has structure at arbitrarily small scales.* If we enlarge part of C repeatedly, we continue to see a complex pattern of points separated by gaps of various sizes. This structure is neverending, like worlds within worlds. In contrast, when we look at a smooth curve or surface under repeated magnification, the picture becomes more and more featureless.
2. *C is self-similar.* It contains smaller copies of itself at all scales. For instance, if we take the left part of C (the part contained in the interval $[0, \frac{1}{3}]$) and enlarge it by a factor of three, we get C back again. Similarly, the parts of C in each of the four intervals of S_2 are geometrically similar to C, except scaled down by a factor of nine.

   If you're having trouble seeing the self-similarity, it may help to think about the sets S_n rather than the mind-boggling set S_∞. Focus on the left half of S_2—it looks just like S_1, except three times smaller. Similarly, the left half of S_3 is S_2, reduced by a factor of three. In general, the left half of S_{n+1} looks like *all* of S_n, scaled down by three. Now set $n = \infty$. The conclusion is that the left half of S_∞ looks like S_∞, scaled down by three, just as we claimed earlier.

   Warning: The strict self-similarity of the Cantor set is found only in the simplest fractals. More general fractals are only approximately self-similar.
3. *The dimension of C is not an integer.* As we'll show in Section 11.3, its dimension is actually $\ln 2/\ln 3 \approx 0.63$! The idea of a noninteger dimension is bewildering at first, but it turns out to be a natural generalization of our intuitive ideas about dimension, and provides a very useful tool for quantifying the structure of fractals.

Two other properties of the Cantor set are worth noting, although they are not fractal properties as such: C *has measure zero* and *it consists of uncountably many points*. These properties are clarified in the examples below.

EXAMPLE 11.2.1:

Show that the ***measure*** of the Cantor set is zero, in the sense that it can be covered by intervals whose total length is arbitrarily small.

Solution: Figure 11.2.1 shows that each set S_n completely covers all the sets that come after it in the construction. Hence the Cantor set $C = S_\infty$ is covered by *each* of the sets S_n. So the total length of the Cantor set must be less than the total length of S_n, for any n. Let L_n denote the length of S_n. Then from Figure 11.2.1 we see that $L_0 = 1$, $L_1 = \frac{2}{3}$, $L_2 = \left(\frac{2}{3}\right)\left(\frac{2}{3}\right) = \left(\frac{2}{3}\right)^2$, and in general, $L_n = \left(\frac{2}{3}\right)^n$. Since $L_n \to 0$ as $n \to \infty$, the Cantor set has a total length of zero. ■

whose base-3 expansion contains no 1's, as claimed. ■

There's still a fussy point to be addressed. What about endpoints like $\frac{1}{3} = .1000\ldots$? It's in the Cantor set, yet it has a 1 in its base-3 expansion. Does this contradict what we said above? No, because this point can also be written solely in terms of 0's and 2's, as follows: $\frac{1}{3} = .1000\ldots = .02222\ldots$. By this trick, each point in the Cantor set can be written such that no 1's appear in its base-3 expansion, as claimed.

Now for the payoff.

EXAMPLE 11.2.3:

Show that the Cantor set is uncountable.

Solution: This is just a rewrite of the Cantor diagonal argument of Example 11.1.4, so we'll be brief. Suppose there were a list $\{c_1, c_2, c_3, \ldots\}$ of all points in C. To show that C is uncountable, we produce a point $\bar{c}$ that is in C but not on the list. Let c_{ij} denote the jth digit in the base-3 expansion of c_i. Define $\bar{c} = \bar{c}_{11}\bar{c}_{22}\ldots$, where the overbar means we switch 0's and 2's: thus $\bar{c}_{nn} = 0$ if $c_{nn} = 2$ and $\bar{c}_{nn} = 2$ if $c_{nn} = 0$. Then $\bar{c}$ is in C, since it's written solely with 0's and 2's, but $\bar{c}$ is not on the list, since it differs from c_n in the nth digit. This contradicts the original assumption that the list is complete. Hence C is uncountable. ■

11.3 Dimension of Self-Similar Fractals

What is the "dimension" of a set of points? For familiar geometric objects, the answer is clear—lines and smooth curves are one-dimensional, planes and smooth surfaces are two-dimensional, solids are three-dimensional, and so on. If forced to give a definition, we could say that *the dimension is the minimum number of coordinates needed to describe every point in the set.* For instance, a smooth curve is one-dimensional because every point on it is determined by one number, the arc length from some fixed reference point on the curve.

But when we try to apply this definition to fractals, we quickly run into paradoxes. Consider the ***von Koch curve***, defined recursively in Figure 11.3.1.

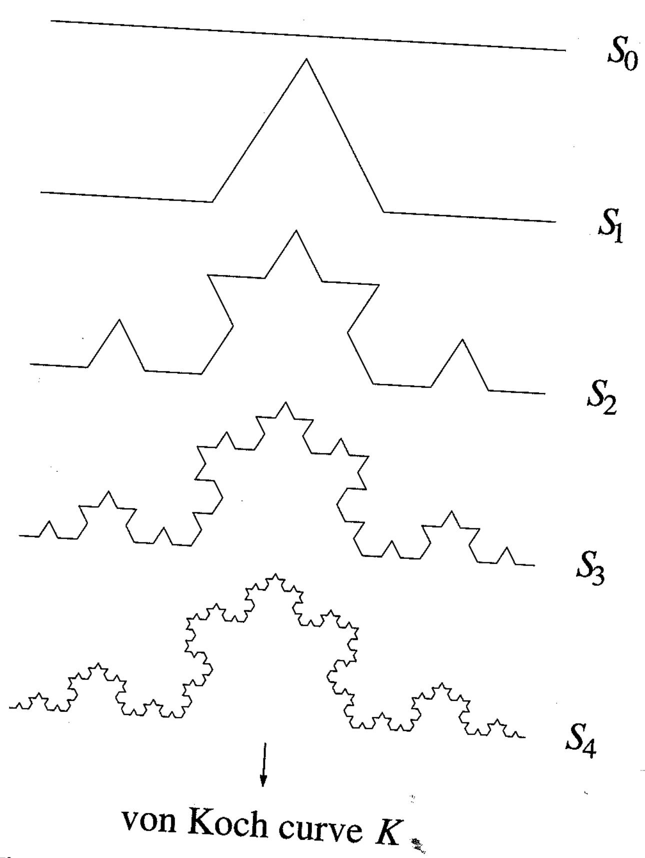

Figure 11.3.1

We start with a line segment S_0. To generate S_1, we delete the middle third of S_0 and replace it with the other two sides of an equilateral triangle. Subsequent stages are generated recursively by the same rule: S_n is obtained by replacing the middle third of each line segment in S_{n-1} by the other two sides of an equilateral triangle. The limiting set $K = S_\infty$ is the von Koch curve.

A Paradox

What is the dimension of the von Koch curve? Since it's a curve, you might be tempted to say it's one-dimensional. But the trouble is that K has *infinite arc length*! To see this, observe that if the length of S_0 is L_0, then the length of S_1 is $L_1 = \frac{4}{3}L_0$, because S_1 contains four segments, each of length $\frac{1}{3}L_0$. The length increases by a factor of $\frac{4}{3}$ at each stage of the construction, so $L_n = (\frac{4}{3})^n L_0 \to \infty$ as $n \to \infty$.

Moreover, the arc length between *any* two points on K is infinite, by similar reasoning. Hence points on K aren't determined by their arc length from a particular point, because every point is infinitely far from every other!

11.4 Box Dimension

To deal with fractals that are not self-similar, we need to generalize our notion of dimension still further. Various definitions have been proposed; see Falconer (1990) for a lucid discussion. All the definitions share the idea of "measurement at a scale ε"—roughly speaking, we measure the set in a way that ignores irregularities of size less than ε, and then study how the measurements vary as $\varepsilon \to 0$.

Definition of Box Dimension

One kind of measurement involves covering the set with boxes of size ε (Figure 11.4.1).

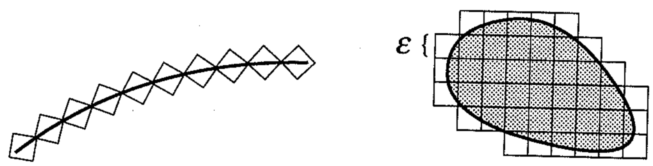

$$N(\varepsilon) \propto \frac{L}{\varepsilon} \qquad\qquad N(\varepsilon) \propto \frac{A}{\varepsilon^2}$$

Figure 11.4.1

Let S be a subset of D-dimensional Euclidean space, and let $N(\varepsilon)$ be the minimum number of D-dimensional cubes of side ε needed to cover S. How does $N(\varepsilon)$ depend on ε? To get some intuition, consider the classical sets shown in Figure 11.4.1. For a smooth curve of length L, $N(\varepsilon) \propto L/\varepsilon$; for a planar region of area A bounded by a smooth curve, $N(\varepsilon) \propto A/\varepsilon^2$. The key observation is that the dimension of the set equals the exponent d in the ***power law*** $N(\varepsilon) \propto 1/\varepsilon^d$.

This power law also holds for most fractal sets S, except that d is no longer an integer. By analogy with the classical case, we interpret d as a dimension, usually called the *capacity* or ***box dimension*** of S. An equivalent definition is

$$d = \lim_{\varepsilon \to 0} \frac{\ln N(\varepsilon)}{\ln(1/\varepsilon)}, \text{ if the limit exists.}$$

EXAMPLE 11.4.1:

Find the box dimension of the Cantor set.

Solution: Recall that the Cantor set is covered by each of the sets S_n used in its construction (Figure 11.2.1). Each S_n consists of 2^n intervals of length $(1/3)^n$, so if we pick $\varepsilon = (1/3)^n$, we need all 2^n of these intervals to cover the Cantor set. Hence

$N = 2^n$ when $\varepsilon = (1/3)^n$. Since $\varepsilon \to 0$ as $n \to \infty$, we find

$$d = \lim_{\varepsilon \to 0} \frac{\ln N(\varepsilon)}{\ln(1/\varepsilon)} = \frac{\ln(2^n)}{\ln(3^n)} = \frac{n \ln 2}{n \ln 3} = \frac{\ln 2}{\ln 3}$$

in agreement with the similarity dimension found in Example 11.3.1. ■

This solution illustrates a helpful trick. We used a discrete sequence $\varepsilon = (1/3)^n$ that tends to zero as $n \to \infty$, even though the definition of box dimension says that we should let $\varepsilon \to 0$ continuously. If $\varepsilon \neq (1/3)^n$, the covering will be slightly wasteful—some boxes hang over the edge of the set—but the limiting value of d is the same.

EXAMPLE 11.4.2:

A fractal that is *not* self-similar is constructed as follows. A square region is divided into nine equal squares, and then one of the small squares is selected at random and discarded. Then the process is repeated on each of the eight remaining small squares, and so on. What is the box dimension of the limiting set?

Solution: Figure 11.4.2 shows the first two stages in a typical realization of this random construction.

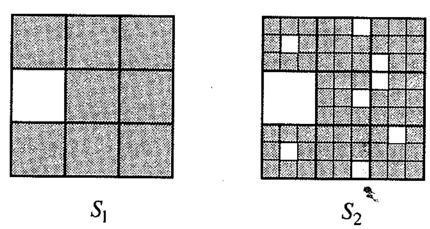

Figure 11.4.2

Pick the unit of length to equal the side of the original square. Then S_1 is covered (with no wastage) by $N = 8$ squares of side $\varepsilon = \frac{1}{3}$. Similarly, S_2 is covered by $N = 8^2$ squares of side $\varepsilon = (\frac{1}{3})^2$. In general, $N = 8^n$ when $\varepsilon = (\frac{1}{3})^n$. Hence

$$d = \lim_{\varepsilon \to 0} \frac{\ln N(\varepsilon)}{\ln(1/\varepsilon)} = \frac{\ln(8^n)}{\ln(3^n)} = \frac{n \ln 8}{n \ln 3} = \frac{\ln 8}{\ln 3}. \ ■$$

Critique of Box Dimension

When computing the box dimension, it is not always easy to find a minimal cover. There's an equivalent way to compute the box dimension that avoids this problem. We cover the set with a square mesh of boxes of side ε, count the number of occupied boxes $N(\varepsilon)$, and then compute d as before.

Even with this improvement, the box dimension is rarely used in practice. Its computation requires too much storage space and computer time, compared to other