



© 1997–2009, Millennium Mathematics Project, University of Cambridge.

Permission is granted to print and copy this page on paper for non-commercial use. For other uses, including electronic redistribution, please contact us.


September 1999

Features



Extracting beauty from chaos

by Andy Burbanks



Extracting beauty from chaos

There are any number of sites on the World Wide Web dedicated to galleries of computer-generated fractal images. Pictures based on Lyapunov Exponent fractals, such as the one pictured above, are some of the most striking and unusual.

Like most fractal images, Lyapunov Exponent fractals are produced by iterating functions and observing the chaotic behaviour that may result.


In this article, we'll be exploring what Lyapunov exponents are, how you iterate functions, what "chaos" really is, and how you go about turning simple maths into marvellous images. There's also a companion article, [Computing the Mandelbrot set](#) by Andrew Williams, which looks at the technical details involved in writing a program to compute the Mandelbrot set.

What is iteration?

Even simple equations can produce complex behaviour when there is feedback in the system. For example, take a simple function like

$$f(x) = 4x(1 - x)$$

The graph of this function is just an upside-down parabola passing through the (x, y) pairs $(0, 0)$, $(1/2, 1)$, $(1, 0)$, as shown in the following figure:



Graph of the function f

The number 4 in the definition of the function is there so that the graph of the function fits neatly into the unit box; in other words, if we apply the function to any number in the unit interval $0 \leq x \leq 1$, then we get another number in the unit interval.

Now, pick a real number $x[0]$ between 0 and 1 (the "initial condition" or "initial point") and apply the function f to this x to get

$$x[1] = f(x[0]) = 4x[0](1 - x[0])$$

What is iteration?

Extracting beauty from chaos

For example, suppose we choose $x[0] = 0.3$; then we get

$$x[1] = f(x[0]) = 4(0.3)(1 - 0.3) = 0.84.$$

Now apply "feedback" to $x[1]$ by applying the function again to get

$$x[2] = f(x[1]) = f(f(x[0]))$$

and keep repeating this process to get

$$x[3] = f(x[2]) = f(f(f(x[0]]));$$

$$x[4] = f(x[3]) = f(f(f(f(x[0]]))),$$

and so on.

So, for each number $n = 0, 1, 2, \dots$, we put

$$x[n+1] = f(x[n]).$$

This process is called "iterating" the function f with the initial condition (or "initial point") $x[0]$. The sequence of numbers

$$x[0], x[1], x[2], x[3], \dots$$

is called the "orbit" of the initial condition $x[0]$.

The butterfly effect

If we add a small error to the initial point $x[0]$, e.g. we look at

$$u[0] = x[0] + 0.001$$

and then we look at the orbit $u[0], u[1], u[2], \dots$ of $u[0]$, we might expect that the small error is not very important.


However, it turns out that a function like the one above tends to *amplify small errors* until they become large, so the difference or "error" between the two orbits, i.e. the sequence

$$u[0] - x[0], u[1] - x[1], u[2] - x[2], \dots, u[k] - x[k], \dots$$

tends to grow until it is as large as the numbers $x[k]$ themselves.


The following diagrams illustrate this phenomenon: first, we show the orbit of an initial point in blue. We plot the points $x[k]$ (for $k = 0, 1, 2$ etc.) against the step-number k : this is called a "time-series".

Extracting beauty from chaos



Orbit of an initial point

In the next picture, we compare this with the orbit of another close-by point, shown in red. In other words, we add a small error and look at what happens to the orbit:

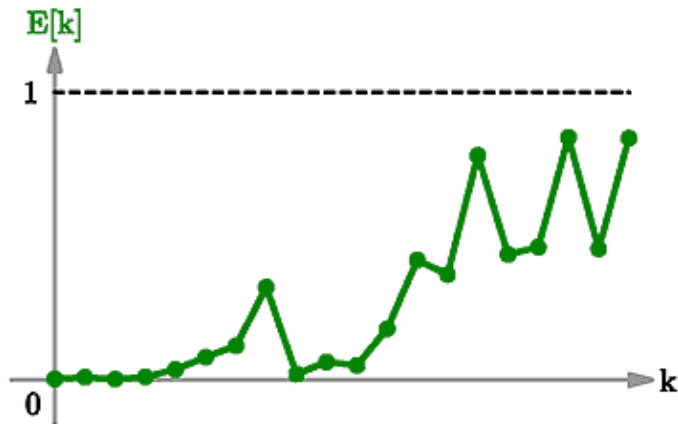


Orbits of two initially close-by points

Notice that for the first few steps (on the left-hand side of the diagram), the orbits remain very close together and it seems that the small error hasn't made much difference. However, after just a few steps (toward the right-hand side of the diagram) the orbits begin to look completely different. The error has grown until the orbits are far apart.

The next picture shows the size of the error (i.e. the difference between the two orbits) at each step:

Extracting beauty from chaos



Error between the orbits

Notice how the error starts off very small, but after a few steps it has grown as large as the orbits themselves. This is like the noise in a radio transmission becoming as large as the signal itself and swamping it completely.

In fact, no matter how tiny we make the initial error, in most cases it will tend to be amplified until it grows large. This behaviour, where small errors tend to become amplified until they are as large as the "signal" itself, is called *sensitive dependence on initial conditions*.

Another popular name for this phenomenon is the "*Butterfly Effect*", so-called because the idea is rather like the tiny flapping of a butterfly's wings becoming amplified until it causes a hurricane on the other side of the world.

This is because the weather is sensitive to small errors in the same way that the function above is, so that – no matter how accurately we try to measure the weather right now (the temperature, humidity, etc.) – inevitably small errors in our measurements are amplified until they make our long-term weather forecasts unreliable.

This effect is *one* of the hallmarks of "chaos": chaotic systems all have this feature (but not all sensitive systems are chaotic!) The other hallmarks of chaos are to do with periodic orbits ("cycles", where there is a repeating pattern in the orbit $x[0], x[1], \dots$) and "mixing" behaviour (like the kneading of dough, where orbits tend to get thoroughly "mixed around" throughout all the possible values).

When all three hallmarks are present (sensitivity, periodic cycles and mixing), we have a phenomenon called "chaos". Unfortunately, Hollywood films have overused the word, so most people think that it means something altogether more wishy-washy, as opposed to a well-defined and precise mathematical term.

Measuring sensitivity: The Lyapunov Exponent

So, we see that when certain systems are iterated, small errors tend to be amplified. How can we measure how quickly this happens?

A quantity called the "Lyapunov Exponent" (after its inventor, [Alexander M. Lyapunov](#), 1857–1918) is used to measure the extent to which ever-smaller "infinitesimal" errors are amplified. In other words, the Lyapunov Exponent measures *how sensitive* a system is to errors.

Growth of small errors

Remember that we begin with a small initial error. Let's call this $E[0]$. So, we begin with


$$u[0] = x[0] + E[0]$$

i.e. initial condition $x[0]$ and error $E[0]$. Then we compare the "exact" orbit $x[0], x[1], x[2], \dots$ against the orbit with the error $u[0], u[1], u[2], \dots$ and see how the difference between them grows.

In other words, we look at the sequence of errors:

$$\begin{aligned} E[0] &= u[0] - x[0], \\ E[1] &= u[1] - x[1], \\ E[2] &= u[2] - x[2], \\ &\dots \end{aligned}$$

The following picture demonstrates how an initial error $E[0]$ grows to give the error $E[1]$ at the next step for the function f described above:



Growth of an initial error

To see how much the error is amplified at each iteration step k , we compute

$$\left| \frac{E[k+1]}{E[k]} \right|$$

i.e. we see how big the next error $E[k+1]$ is, in comparison with the current error $E[k]$. (For example, if

Extracting beauty from chaos

the next error were twice as big as the current one, then the expression above would have the value 2.)

After n steps, the initial error $E[0]$ has been amplified by a factor of:

$$\left| \frac{E[n]}{E[0]} \right| = \left| \frac{E[n]}{E[n-1]} \right| \cdot \left| \frac{E[n-1]}{E[n-2]} \right| \cdots \left| \frac{E[1]}{E[0]} \right|$$

Now, suppose we had instead just looked at the simple linear function

$$g(x) = cx,$$

where c is some constant greater than 1. Then any error is simply magnified by c at each iteration:

$$\begin{aligned} g(x + E) &= c(x + E) \\ &= cx + cE \\ &= g(x) + cE; \\ g(g(x) + cE) &= c(g(x) + cE) \\ &= cg(x) + (c^2)E \\ &= g(g(x)) + (c^2)E, \end{aligned}$$

so that after n steps, the initial error has been magnified by

$$\left| \frac{E[n]}{E[0]} \right| = c^n$$

(Note that if c were less than 1, the error would actually be decreased rather than increased.)

If we want to find out what the value of c is for such a linear function g , we have

$$\ln(c) = \frac{1}{n} \ln \left(\left| \frac{E[n]}{E[0]} \right| \right)$$

This is the idea behind the Lyapunov exponent: we take a *general* function f and use this same formula to see how much small errors tend to be magnified. We take more and more iterates (larger and larger n) and calculate the corresponding values for c . If, as we let n grow larger and larger, these values settle down to a constant, then this suggests that the error tends to grow on average like c^n .

Taking the logarithm turns a product into a sum:

$$\frac{1}{n} \ln \left(\left| \frac{E[n]}{E[0]} \right| \right) = \frac{1}{n} \ln \left(\left| \frac{E[n]}{E[n-1]} \right| \cdot \left| \frac{E[n-1]}{E[n-2]} \right| \cdots \left| \frac{E[1]}{E[0]} \right| \right)$$

Extracting beauty from chaos

$$= \left(\frac{1}{n}\right) \sum_{k=1}^{k=n} \ln \left(\left| \frac{E[k]}{E[k-1]} \right| \right)$$

The Lyapunov Exponent is defined to be the limiting value of the above quantity as the initial error $E[0]$ is made ever smaller—and-smaller, and the number of iterations n is sent to infinity.

Lyapunov exponent of a differentiable function

Notice that

$$E[k] = f(x[k-1] + E[k-1]) - f(x[k-1])$$

so that

$$\frac{E[k]}{E[k-1]} = \frac{f(x[k-1] + E[k-1]) - f(x[k-1])}{E[k-1]}$$

Now, suppose that the function f is smooth (i.e. we can differentiate it to get $f'(x)$). In this case, you should notice that the expression on the right-hand side looks familiar! Let's rewrite it like this (put $h = E[k-1]$ and $x = x[k-1]$):

$$\frac{E[k]}{E[k-1]} = \frac{f(x+h) - f(x)}{h}$$

If we let h go to zero, then this is just the expression for the derivative! So, as h tends to zero we have:

$$\frac{E[k]}{E[k-1]} \rightarrow f'(x[k-1])$$

Thus we have

$$\begin{aligned} \left(\frac{1}{n}\right) \ln \left(\left| \frac{E[n]}{E[0]} \right| \right) &= \left(\frac{1}{n}\right) \sum_{k=1}^{k=n} \ln \left(\left| \frac{E[k]}{E[k-1]} \right| \right) \\ &= \frac{1}{n} \sum_{k=1}^{k=n} \ln (| f'(x[k-1]) |) \end{aligned}$$

Order and Chaos in the same system

Even some very simple systems are capable of showing *both* orderly and chaotic behaviour.

These systems may make transitions (called "bifurcations") from one type of behaviour to another, as some parameter is varied, rather like tuning-in to different radio stations by turning the dial on a radio.

For example, remember our function

Lyapunov exponent of a differentiable function


Extracting beauty from chaos

$$f(x) = 4x(1-x)$$

whose graph was a parabola through $(x, y) = (0, 0), (1/2, 1),$ and $(1, 0)$. We can add a parameter P to let us move the height of the parabola up and down:

$$f_p(x) = (4p)x(1-x)$$

For a fixed number P (with magnitude between 0 and 1) the graph of f_p against x is a parabola still passing through $(0, 0)$ and $(1, 0)$, but now having its "top" (turning point) at $(1/2, p)$, as shown in the following picture:



Graph of function with a parameter p

You can think of the parameter P as being like the radio tuning dial: adjusting the value of P moves the top of the parabola up and down smoothly, but as we will soon see the behaviour of the function can change suddenly. When tuning a radio, you sometimes happen to hit the right wavelength and hear a different radio station. In the same way, when the value of P (the tuning dial) passes certain special values, the behaviour of orbits can change suddenly.

There is a web page where you can experiment interactively with changing the value of the parabola and the position of the initial point $x[0]$:


see [Further Reading](#) for details.

Plotting the stability

One idea to help us understand the way that the behaviour changes (as we change the parameter) is to plot the value of the Lyapunov exponent (for some chosen initial point $x[0]$) against the parameter P . This will let us estimate how the stability of our system changes as P changes.

The following diagram shows a plot of the calculated Lyapunov exponent (using the starting point $x[0] = 1/2$) against the parameter P .

Extracting beauty from chaos



Lyapunov exponent vs. parameter

Notice the large negative spikes pointing downward. In reality, these spikes in the graph go off to minus infinity, showing that the behaviour of the function is incredibly stable at those parameter values. Remember that we have calculated an estimate of the *logarithm* of the error-multiplier: if this value is very large and negative, then the error-multiplier is very small, which indicates that errors tend to vanish away and the orbit of the system is super-stable.

Between the spikes, the graph touches the p -axis (so that the Lyapunov exponent is zero): this indicates that these parameters give an orbit which is on the edge of stable and unstable behaviour. It is here that sudden changes (bifurcations) in the behaviour of orbits take place.

Notice that the graph stays below zero until it reaches a point near the right hand side of the plot: this is where sensitivity first appears. As the graph rises above the axis, this means that the Lyapunov exponent is positive so that errors tend to be magnified. Even so, in this right-hand region there are still spikes where super-stable orbits appear.

There is much complicated structure in this diagram and even a simple one-parameter family of functions (like our parabola functions) still holds mysteries for mathematical research. There are also features of this plot that are universal, that is to say that we see them even when we look at very complicated functions, rather than our simple example; looking at simple situations in mathematics can sometimes throw light on more complicated ones.

Lyapunov-exponent pictures

To make a two-dimensional picture, we look at functions with *two* parameters (say p and q), rather than just one. Then we use a colour to represent the Lyapunov exponent for each choice of the parameters (p, q) in some rectangle.

Conveniently, we can put these two-parameter functions together from a pair of one-parameter functions. For example, we can simply use the parabola function above. For numbers p and q we consider the functions:

Extracting beauty from chaos

$$P(x) = (4p)x(1-x)$$

$$Q(x) = (4q)x(1-x)$$

Then, to make a single function, we combine them in some way. One way to do this is to apply them, one after the other, in some order. For example, we might apply P first, then Q . So we would make our function f from:

$$f(x) = Q(P(x))$$

(i.e. we are alternating between P and Q) and we then iterate it as usual. (This means that, for example, $f(f(f(x))) = Q(P(Q(P(Q(P(x))))))$.) Or, we might choose

$$f(x) = Q(P(P(P(Q(P(Q(x)))))))$$

or any other sequence of the functions P and Q .

Once we have decided what sequence of P and Q to use – let's choose $Q(P(x))$ – we keep it fixed from now on.

Next, we work out the Lyapunov exponent for parameters (p, q) in a grid. We're going to build our picture one dot (pixel) at a time.

For each (p, q) we take the initial point $x[0] = 1/2$ and iterate f to estimate the Lyapunov exponent, using the Chain Rule of differentiation to see that

$$f'(x) = (Q'(P(x))) \cdot P'(x)$$


and so we calculate the numbers

$$\ln |f'(x[k-1])|$$

for k from 1 up to some n , and average them. This gives us an estimate of the Lyapunov Exponent. (Choosing a larger value for n generally gives a better estimate, but choosing n too large means that a computer program will take a very long time to do the calculation for each point (p, q) .)


We then colour the point corresponding to (p, q) in our image, depending on the value that we estimated for the Lyapunov exponent. For example, we might choose dark reds for positive Lyapunov Exponents and Yellows for negative values, etc. We can choose a smoothly-spread range of colours when producing our pictures.

Extracting beauty from chaos



Part of the Lyapunov exponent picture for the QP sequence of maps

Understanding the pictures



Another Lyapunov Exponent picture

The sharply-defined curves which sometimes appear in these pictures correspond to very large negative Lyapunov exponents (meaning that the system is very stable, in fact "super"-stable, and errors tend to vanish away).

The background regions indicate positive exponents (where the system is sensitive, i.e. small errors tend to grow, and may be chaotic).

The presence of regions that appear to "cross over" each other reveals that there are two or more different

Extracting beauty from chaos

types of behaviour that are "competing" with each other (in fact, these are competing "attractors") – different behaviours "win" the competition for different values of the parameters p and q , so we see what look like different shapes crossing over each other.

There are also repeating patterns and progressions of shapes, revealing lots of structure that would be difficult to guess just from looking at the simple equations used to generate the pictures. By visualising such structures and seeking to understand them, we can make progress in understanding the behaviour of more complicated models of systems in the real world. The beauty of these, and other, fractal images is an invitation to try and understand such patterns.

Further Reading

Iteration of the logistic map:

If your browser can handle JAVA programs, then you can play with changing the height of the parabola and the position of the initial point $x[0]$ (see [Order and Chaos in the same system](#)) by looking at the web page [Iteration of the Logistic Map](#).

Companion Article

- [Computing the Mandelbrot set](#) in this issue.

Other PASS Maths articles on fractals:

- [The origins of fractals in Issue 6](#).
 - [Modelling nature with fractals in Issue 6](#).
 - [Jackson's fractals](#) in this issue.
-

About the author



[Andy Burbanks](#) took a BSc in Mathematics and Computation at the University of Loughborough, and stayed on to do a PhD in Mathematics, during which time he developed an interest in mathematical artwork as a hobby.

He then spent a year as a mathematical artist, working with Bristol-based sculptor Simon Thomas (see [The art of numbers in Issue 8](#)).

He is currently doing research in dynamical systems in the [Pure Maths Department](#) at [Cambridge University](#).

In addition to his research work, Andy is currently designing mathematical posters for the [World Mathematical Year 2000's Posters in the London Underground](#) project.



Plus is part of the family of activities in the Millennium Mathematics Project, which also includes the NRICH and MOTIVATE sites.