The Curious Order of Chaos

Raphael Vicas

Our feeling for beauty is inspired by the harmonious arrangement of order and disorder.

Gert Eilenberger

Section 1: How a Butterfly Ruined Weather Forecasting

Beauty rarely belongs to perfect symmetry. What makes the world compelling is the interplay between order and disorder: patterns that hold, but never quite repeat. Chaos theory gives it a language.

Edward Lorenz, a meteorologist at MIT, discovered chaos by accident in the early 1960s. He was studying weather patterns, a problem of enormous complexity, using a deceptively simple model of air convection. It was built from three coupled non-linear differential equations:

$$\frac{dx}{dt} = \sigma(y - x), \quad \frac{dy}{dt} = x(\rho - z) - y, \quad \frac{dz}{dt} = xy - \beta z$$

The variables x, y, z represented elements of fluid motion (velocity, temperature, and convection intensity, respectively) while the parameters σ, ρ, β controlled the system's behaviour.

One afternoon, Lorenz restarted a computation using slightly rounded values (0.506 instead of 0.506127), resulting in a difference so small it seemed irrelevant. Letting the computer do its magic, he went to get a coffee, as a task that large would take a long, long time. When the results printed out, they diverged completely *(fig.1)*. The trajectories no longer matched. That rounding error, microscopic in scale, had grown into an entirely different weather pattern.

He called it the *butterfly effect*: a butterfly flapping its wings in London in 1342 might contribute to a hurricane in Texas in 1961. It seems almost absurd, and yet the mathematics behind it is sound. Lorenz had stumbled upon chaos - not randomness, but deterministic unpredictability. The system's rules hadn't changed, only the starting point. Plotting the solutions produced an image now famous: the *Lorenz attractor* (fig. 2). A shape looping endlessly, coincidentally like the twin wings of a butterfly, each trajectory bounded yet never repeating. The curves twist around one another, never touching, never escaping.

It is a common misconception that the flapping of a butterfly's wings is entirely the *cause* of a hurricane; this is not the case. The flap shifts a point, A, off its trajectory on the attractor, but then it finds its way back onto it, starting as a different point, B. It traces a similar shape, but at a different time, at a different point in space. So the butterfly isn't powerful enough to create the hurricane, only to steal it. To move it through time, like sliding a bookmark along a page. The hurricane was always coming. The butterfly just decided where and when. However, if we were to have a super-butterfly, whose flaps shake the earth beneath our feet, the gargantuan flap could alter the system enough to move it onto an entirely different attractor. In reality, we don't have super-butterflies, but we do have climate change.

The sensitivity Lorenz found can be described formally using the *Lyapunov exponent*, a measure of how quickly nearby trajectories diverge over time.

Let the initial condition be x_0 and a nearby point be $x_0 + \delta_0$.

Let δ_n denote the separation of the corresponding orbits after n iterations.

Let $f : \mathbb{R}^n \to \mathbb{R}^n$ be differentiable, $x_0 \in \mathbb{R}$, and define the orbit $x_i = f^i(x_0)$.

If $|\delta_n| \approx |\delta_0|\, e^{n\lambda}$, then λ is called the Lyapunov exponent.

Precisely: $\delta_n = f^n(x_0 + \delta_0) - f^n(x_0)$,

$$\Rightarrow \lambda \approx \frac{1}{n}\ln\left|\frac{\delta_n}{\delta_0}\right| = \frac{1}{n}\ln\left|\frac{f^n(x_0+\delta_0) - f^n(x_0)}{\delta_0}\right|$$

$$\Rightarrow \lambda \approx \frac{1}{n}\ln|(f^n)'(x_0)|$$

Additionally, we can show that by the chain rule,

$$(f^n)'(x_0) = (f \circ f^{n-1})'(x_0) = f'(f^{n-1}(x_0)) \cdot (f^{n-1})'(x_0).$$

Applying the same argument recursively gives

$$(f^n)'(x_0) = f'(x_{n-1})\, f'(x_{n-2}) \cdots f'(x_0) = \prod_{i=0}^{n-1} f'(x_i).$$

Therefore,

$$\lambda \approx \frac{1}{n}\ln\left(\prod_{i=0}^{n-1} |f'(x_i)|\right)$$

$$= \frac{1}{n}\sum_{i=0}^{n-1} \ln|f'(x_i)|$$

Therefore, the Lyapunov exponent is: $\lambda = \lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} \ln |f'(x_i)|$.

We can see that the Lyapunov exponent measures the average rate of separation or contraction of nearby trajectories as they evolve throughout the state space. If $\lambda > 0$, nearby trajectories diverge exponentially, making long-term prediction impossible. To calculate λ, one can linearise the system along a trajectory using the *Jacobian* matrix J.

Consider the Lorenz system:

$$\frac{dx}{dt} = \sigma(y - x),$$
$$\frac{dy}{dt} = x(\rho - z) - y,$$
$$\frac{dz}{dt} = xy - \beta z,$$

To study how small perturbations $\delta\mathbf{x} = (\delta x, \delta y, \delta z)^T$ evolve, we linearise the system around a trajectory. Using a first-order Taylor expansion:

$$\frac{d}{dt}(\mathbf{x} + \delta\mathbf{x}) = \mathbf{F}(\mathbf{x} + \delta\mathbf{x}) \approx \mathbf{F}(\mathbf{x}) + J(\mathbf{x})\delta\mathbf{x},$$

where $J(\mathbf{x})$ is the Jacobian matrix defined as

$$J_{ij} = \frac{\partial F_i}{\partial x_j}.$$

The Lorenz vector field is:

$$\mathbf{F}(x, y, z) = \begin{pmatrix} \sigma(y - x) \\ x(\rho - z) - y \\ xy - \beta z \end{pmatrix}.$$

Compute each partial derivative:

$$\frac{\partial}{\partial x}[\sigma(y - x)] = -\sigma, \quad \frac{\partial}{\partial y}[\sigma(y - x)] = \sigma, \quad \frac{\partial}{\partial z}[\sigma(y - x)] = 0,$$
$$\frac{\partial}{\partial x}[x(\rho - z) - y] = \rho - z, \quad \frac{\partial}{\partial y}[x(\rho - z) - y] = -1, \quad \frac{\partial}{\partial z}[x(\rho - z) - y] = -x,$$
$$\frac{\partial}{\partial x}[xy - \beta z] = y, \quad \frac{\partial}{\partial y}[xy - \beta z] = x, \quad \frac{\partial}{\partial z}[xy - \beta z] = -\beta.$$

Combine the partial derivatives into a 3×3 matrix:

$$J(x, y, z) = \begin{bmatrix} -\sigma & \sigma & 0 \\ \rho - z & -1 & -x \\ y & x & -\beta \end{bmatrix}.$$

Sustained stretching along trajectories produces a positive largest Lyapunov exponent, which is the basis for chaos. The Jacobian thus provides a linearised approximation of the local stretching and contraction in the phase space.

Applying this to the Lorenz attractor:
First, track a trajectory $(x(t), y(t), z(t))$. Next, each point, compute $J(x, y, z)$. The eigenvalues of J describe how nearby points diverge, which explains the 'butterfly effect' and the sensitive dependence on initial conditions.

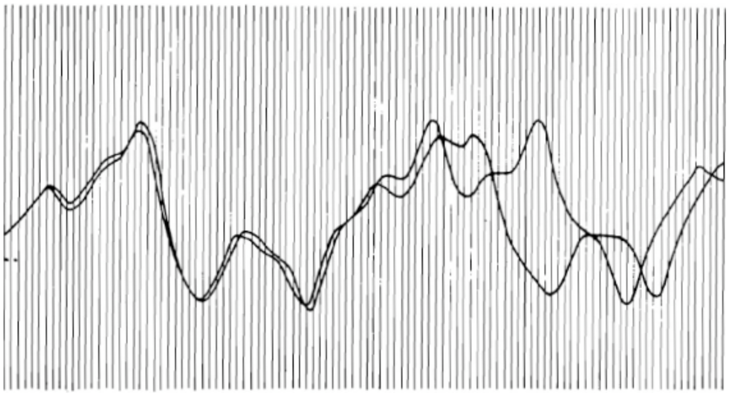

Figure 1: *Starting from almost the same position, you can see the two trajectories growing further apart, becoming incoherent.*

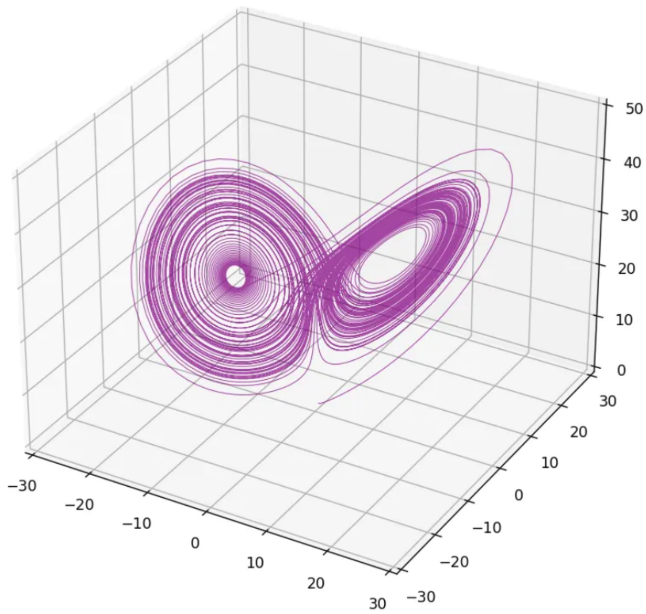

Figure 2: *The Lorenz attractor plotted static in phase space.*

Chaos is highly prevalent in nature. Whether it be in plants or in the weather. Perhaps beauty isn't the triumph of order over disorder, but their uneasy friendship. Too much symmetry and it's boring. Too much noise, and we get lost. Somewhere in between that fragile balance, life begins to sing.

Benoît Mandelbrot called this '*the geometry of nature*.' Classical geometry fails to capture reality. But fractals and chaos do. They explain a hidden order behind apparent disorder. Imperfection can be measurable, and even beautiful. Equations aren't cold abstractions; they breathe and carry tension. I think that's what Eilenberger's quote captures. Beauty isn't perfection, but the tension between holding together and falling apart. The universe is a conversation, never perfectly repeatable, never entirely silent. It is alive because it is never quite still.

Smale's horseshoe map (fig.3) offers another angle. Imagine a square of dough. You stretch it horizontally so it becomes long and thin, compress it vertically so it becomes narrow, then bend it into a horseshoe shape and lay it back over the original square. Iterate this process again and again. Were you to have placed a point (or perhaps a chocolate chip) on the square, you would notice that predicting where the point appears after each stretch and fold becomes increasingly difficult.

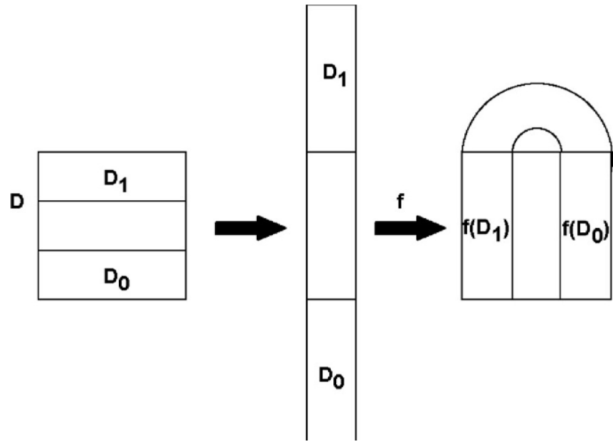

Figure 3: *Smale's horseshoe map: where a region of phase space is stretched, folded, and reinserted, producing an invariant set with sensitive dependence on initial conditions.*

We can see how predicting where a point will end up becomes increasingly difficult through computation. Consider a simple system of two coupled differential equations:

$$\frac{dx}{dt} = x - y, \quad \frac{dy}{dt} = x + y - x^2 y$$

The system can be integrated numerically using the *fourth-order Runge–Kutta method* (RK4), a technique that evaluates intermediate slopes to approximate the solution of differential equations with high accuracy. Its global error scales as Δt^4. Since chaotic systems cannot be solved analytically in most cases, we rely on numerical methods such as RK4 to trace their evolution.

Starting with initial conditions $(x_0, y_0) = (0.1, 0.1)$ and integrating using RK4:

$$
\begin{aligned}
k_1^{x,y} &= f(x_n, y_n) \\
k_2^{x,y} &= f(x_n + k_1^x \frac{\Delta t}{2}, y_n + k_1^y \frac{\Delta t}{2}) \\
k_3^{x,y} &= f(x_n + k_2^x \frac{\Delta t}{2}, y_n + k_2^y \frac{\Delta t}{2}) \\
k_4^{x,y} &= f(x_n + k_3^x \Delta t, y_n + k_3^y \Delta t)
\end{aligned}
$$

We can write this system in vector form as $\frac{d\mathbf{y}}{dt} = f(\mathbf{y})$, where

$$
f(x, y) = \begin{pmatrix} x - y \\ x + y - x^2 y \end{pmatrix}.
$$

The next step is:

$$
x_{n+1} = x_n + \frac{\Delta t}{6} \left(k_1^x + 2k_2^x + 2k_3^x + k_4^x \right)
$$

$$
y_{n+1} = y_n + \frac{\Delta t}{6} \left(k_1^y + 2k_2^y + 2k_3^y + k_4^y \right)
$$

For our system and our chosen values for x_0 and y_0, the RK4 computes:

$$
(x_1, y_1) \approx (0.09999, 0.1019998)
$$

with a chosen time step $\Delta t = 0.01 \approx \frac{y_1 - y_0}{k_1^y}$

After just a few dozen steps, tiny differences in initial conditions, say, $x_0 = 0.1001$ instead of 0.1, produce trajectories that diverge relatively dramatically ($x_1 = 0.1000927$). Plotting y versus x reveals looping, bounded curves: strange attractors.

Once you notice chaos, it starts appearing everywhere. Turbulent currents in a flowing river, clouds dissipating against the sky, and a heartbeat are all bound but never identical. Each carries rules, yet evolves within them. Artists have always sensed this. Musicians stretch a rhythm just a fraction, enough for tension. Painters let a line wander. Writers sometimes let a sentence collapse under its own weight. It's not simply carelessness, but the very pulse of life. Craft is about negotiating that tension and about letting structure and freedom meet. Humans respond to this hidden order instinctively. Just as Lorenz, Smale, and Mandelbrot uncovered patterns in abstract systems, we seek them in the world around us. Chaos is not alien. It is very much human. It is a lens through which complexity and beauty can be understood.

Section 2: Bad Skiing or Chaos?

Chaos isn't limited to atmospheric models; it applies to anything you can think of. Take my favourite sport, skiing, as an example. A mogul-covered slope looks simple enough at first. You might think it's just a series of bumps, something skiers navigate while breaking their backs and knees. But step closer, and it's buzzing with complexity. Every subtle bump in the snow nudges a trajectory in unexpected ways (a good excuse for falling).

Mathematically, we can describe the slope with:

$$H(x, y) = -ax - b\cos(px)\cos(qy)$$

where x measures the downhill direction, y the cross-slope, a the average incline, b the bump height, and the parameters p and q control how frequently those bumps appear. Neat on paper, but then a skier enters, and neatness begins to fray.

The motion of a ski or board sliding along this surface is governed by

$$\frac{dx}{dt} = u, \quad \frac{dy}{dt} = v,$$

$$\frac{du}{dt} = -\frac{\partial H}{\partial x} - c(H)u, \quad \frac{dv}{dt} = -\frac{\partial H}{\partial y} - c(H)v$$

where u and v denote downhill and cross-slope velocities. The damping coefficient $c(H)$ represents friction and depends on the local slope of the surface. If $H_x = \partial H/\partial x$ and $H_y = \partial H/\partial y$, one simple model is

$$c(H) = c_0\sqrt{1 + H_x^2 + H_y^2}.$$

This reflects the fact that the normal force (and therefore friction) changes as the steepness of the terrain varies. The partial derivatives of the slope's height are:

$$\frac{\partial H}{\partial x} = -a + bp\sin(px)\cos(qy), \quad \frac{\partial H}{\partial y} = bq\cos(px)\sin(qy)$$

Even more structure appears when the motion is viewed as a non-linear oscillator. The cross–slope dynamics are governed by the force $-\frac{\partial H}{\partial y} = -bq\cos(px)\sin(qy)$. Near the centre of a valley, where y is small, we can approximate $\sin(qy) \approx qy$, giving a restoring force approximately proportional to $-y$, as $-\frac{\partial H}{\partial y} \approx -bq^2\cos(px)y$. In this system, the skier behaves like a harmonic oscillator moving across the slope. As the skier travels downhill with velocity u, the periodic bumps introduce a forced frequency roughly pu ($x \approx ut$). For small bump heights b the forcing is weak and the motion remains regular, producing smooth oscillations across the slope. As b increases, however, resonances develop between the natural oscillation of the skier and the periodic forcing of the terrain. These resonances destabilise the regular motion, leading to period–doubling and eventually chaotic trajectories.

To explore the sensitive dependence on initial conditions, we integrate the system numerically using the RK4 method once again. Suppose two boards start almost identically:

Board 1: $x = 0, y = 0, u = 3.5, v = 0$

Board 2: $x = 0, y = 0.001, u = 3.5, v = 0$

After integrating using RK4, for a moment, they move almost in sync, brushing past the same bumps. Then slowly, inevitably, they part ways, carving divergent paths. Sometimes they cross again. When you plot v versus y, a strange attractor emerges. The slope sets the frame, then lets the story unfold.

Exploring bifurcations adds another layer of insight. Gradually increasing the bump height b transforms the slope's dynamics:

$$\begin{aligned} b = 0 &\rightarrow \text{flat slope} \rightarrow \text{linear, predictable motion} \\ b = 0.11 &\rightarrow \text{slight oscillations, still periodic} \\ b = 0.4 &\rightarrow \text{first chaotic trajectories appear, points scatter in phase space} \\ b = 0.5\text{–}0.6 &\rightarrow \text{fully chaotic, multiple velocities coexist unpredictably} \end{aligned}$$

By plotting the local maxima of v for each b, we construct a bifurcation diagram *(fig.4)*. Lines represent periodic motion, clouds indicate chaos. Just as the ski slope folds order into chaos, so do small perturbations in relationships, ecosystems, and human decisions produce complexity, beauty, and surprise.

Skiing such a slope isn't always reckless. It can be a dance where gravity insists, friction argues, and bumps interrupt. The skier must adapt and sense the terrain. Each turn is unique. The slope frames the possibility but doesn't dictate exact movement. You feel it in your muscles, in the snow's crunch under your skis, in the wind brushing past your face. The equations exist, but experience is intimate.

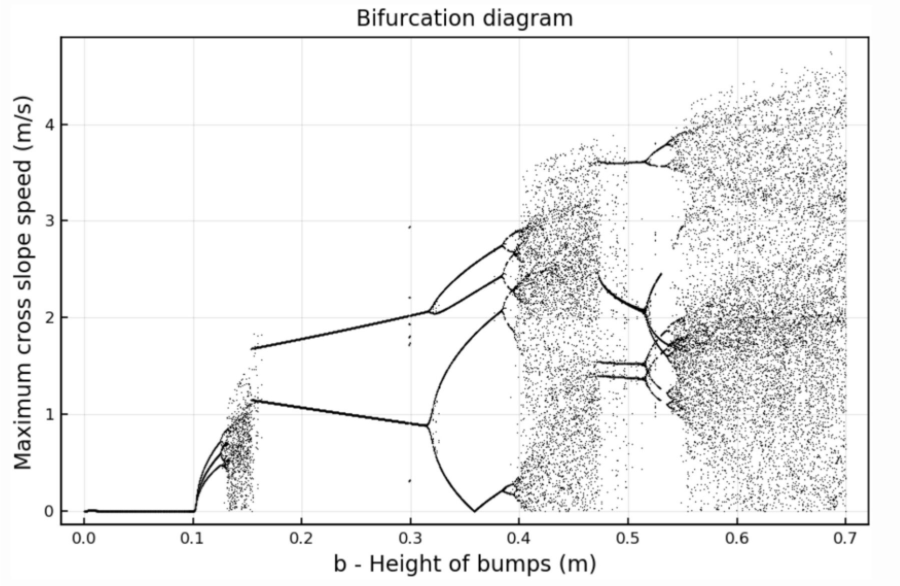

Figure 4: *Bifurcation diagram of maximum cross slope speed against height of mogul bumps.*

What happens when the slope itself can retain memory of motion, one may ask? Each pass of a skier slightly compresses or disturbs the snow, subtly altering friction. We model this with a memory-dependent friction coefficient:

$$c(t) = c_0 + \beta\sqrt{1 + H_x^2 + H_y^2} + \alpha \int_{t-\tau}^{t} \left(u(s)^2 + v(s)^2\right) dt$$

Here, c_0 is the base friction, and α and β are scaling constants. The second term accounts for the local geometry of the slope. As the surface becomes steeper, the normal force and effective contact increase, and friction grows accordingly.

The final term introduces memory. Rather than depending only on the current state, the friction now reflects the recent history of motion. The integral measures the intensity of movement over the past τ seconds, so faster or more erratic motion leaves a lingering effect. In this sense, the snow does not simply respond instantaneously, but retains a trace of what has just occurred.

Small differences in initial conditions are still amplified along trajectories, a feature captured by the largest Lyapunov exponent, $\lambda_{\max}$. However, the presence of memory modifies this behaviour. Numerical exploration suggests that short memory can slightly reduce divergence by smoothing rapid fluctuations, while intermediate memory can enhance it by reinforcing variations in motion.

For longer memory, an additional timescale is introduced into the system. This interacts with the natural oscillations of the skier across the slope, and the motion no longer settles into a simple repeating pattern. Instead, it becomes *quasi-periodic*: when a system is influenced by multiple underlying cycles that do not share a common period. The result is a pattern that appears ordered and repeating at first glance, yet evolves without ever closing into a true cycle. Trajectories fold and twist through phase space, reflecting an interaction between past and present motion. When combined with small perturbations or noise, this produces intricate, evolving patterns and dynamically shifting attractors. By using a *non-Markovian* system, we unravel a whole tapestry of motion where both past and present are entangled. This tapestry, it turns out, has a texture. Look closely enough at the shapes chaos leaves behind, and something unexpected appears: structure, repeating, all the way down.

Section 3: Just One More Zoom, I Promise

Fractals appear throughout nature: in coastlines, lightning, and leaf veins. At first, they seem irregular, yet closer inspection reveals repeating structure at every scale. Benoît Mandelbrot recognised these patterns as a new kind of geometry, one capable of describing the roughness of the natural world.

A bit of preamble before getting into the fun stuff: A fractal is a pattern that carries traces of itself at every scale. Look closely, and smaller pieces resemble the larger shape, though not always perfectly. Some fractals are exact and orderly, built from precise rules. Others are rougher, the kind you find in nature, where the repetition can only be approximate. Nonetheless, they both share the same idea: the deeper you look, the more structure appears, as if the pattern keeps unfolding just beyond the point where you expect it to stop.

Take the *Koch snowflake (fig.5)*. You start with a triangle. On each side, one-third of the way along, add a smaller triangle. On each new triangle, repeat this process.

$$P_n = P_0 \left(\frac{4}{3}\right)^n$$

while the area converges:

$$A_n = A_0 + \sum_{k=1}^{n} \frac{1}{9^k} A_0$$

The perimeter grows endlessly, yet the area barely moves. A more interesting version of this is the *Menger Sponge (fig.6)*, whose surface area diverges to infinity, yet the volume converges to zero. Fractals such as these, however, do not fully encapsulate what it means to be fractal. Just from encountering these forms of fractals, you would assume they are all uniform with no variance, but don't let this nicety distract you from the truth; fractals are usually erratic and rough, with infinite beauty to explore. Such a simplification is much like saying a penguin is spherical, or a human is a cuboid; you've taken everything unique away from it!

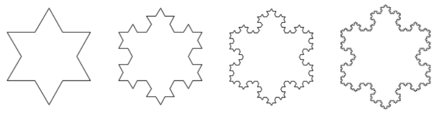

Figure 5: *First few iterations to generate the Koch snowflake*

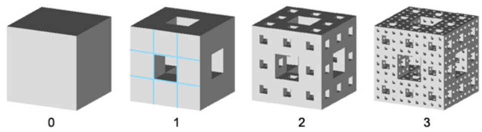

Figure 6: *First few iterations to generate the Menger sponge*

If we now consider the *logistic map*:

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

It originated as a model for population growth. Here, x_n represents the fraction of the maximum population at generation n, and r is the growth rate. At low values of r, the system behaves predictably. The population tends towards a single stable value, independent of small differences in the starting point. The dynamics are simple, and the map feels entirely tame.

To understand when the behaviour of the logistic map changes, we examine its fixed points. A fixed point satisfies

$$x^* = rx^*(1 - x^*).$$

Solving gives two fixed points:

$$x_1^* = 0, \qquad x_2^* = 1 - \frac{1}{r}.$$

To determine whether these points are stable, we examine the derivative

$$f'(x) = r(1 - 2x).$$

A fixed point is stable when small perturbations shrink over time, which occurs when

$$|f'(x^*)| < 1.$$

Evaluating this at x_2^*:

$$f'(x^*) = r\left(1 - 2\left(1 - \frac{1}{r}\right)\right) = 2 - r.$$

Thus stability requires

$$|2 - r| < 1,$$

which gives

$$1 < r < 3.$$

At $r = 3$, the fixed point loses stability and the system undergoes its first *period–doubling bifurcation*. As r increases, the system begins to bifurcate further. A single fixed point becomes unstable, and those two values split again, producing four, then eight, sixteen. Each bifurcation occurs faster than the last, accumulating at a critical point $r_\infty \approx 3.56995$. Beyond this threshold, the map enters chaos: trajectories no longer settle into fixed points or repeating cycles. Tiny changes in initial conditions grow exponentially, making long-term prediction impossible.

Mitchell Feigenbaum studied these period-doubling cascades in detail. He discovered a remarkable constant, now known as the *Feigenbaum* number $\delta \approx 4.669$. This value describes the ratio of intervals between successive bifurcations along the r-axis *(fig.7)*:

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

What makes this constant extraordinary is its *universality*. It is not specific to the logistic map. Any one-dimensional map with a single quadratic maximum exhibits the same period-doubling ratio. Chemical reactions, electrical circuits, fluid flows, and population models all bifurcate in the same rhythm, converging on the same number. Feigenbaum discovered this with a pocket calculator in the 1970s, initially not believed by the journals to which he submitted. The universe, it seems, has a favourite way to tip into chaos, and it always does so at the same rate.

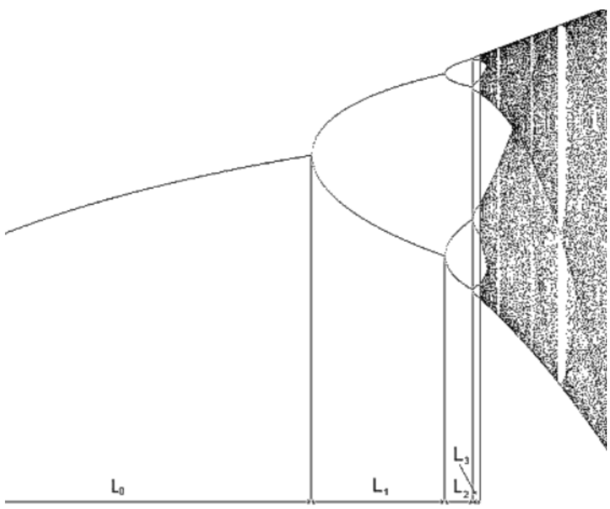

Figure 7: The Feigenbaum value can be calculated by the ratio of $L_n - L_{n-1} : L_{n+1} - L_n$

The *Mandelbrot set (fig.8)* extends these ideas into the complex plane. Consider:

$$z_{n+1} = z_n^2 + c.$$

Here, z and c are complex numbers. The Mandelbrot set is the collection of c values for which the sequence remains bounded. Zooming into the boundary reveals an infinite nesting of structure. Each 'bulb' of the set corresponds to a region of parameter space where a specific period dominates. The main cardioid corresponds to stable fixed points, secondary bulbs to period-two, period-four cycles, and so on. Each bulb contains its own miniature bifurcation structure *(fig.9)*, *(fig.10)*, effectively a logistic map in miniature. The self-similarity of the Mandelbrot set demonstrates the same universality seen in Feigenbaum's work: the ratios of bifurcations, the structure of period doubling, and the onset of chaos all repeat at every scale.

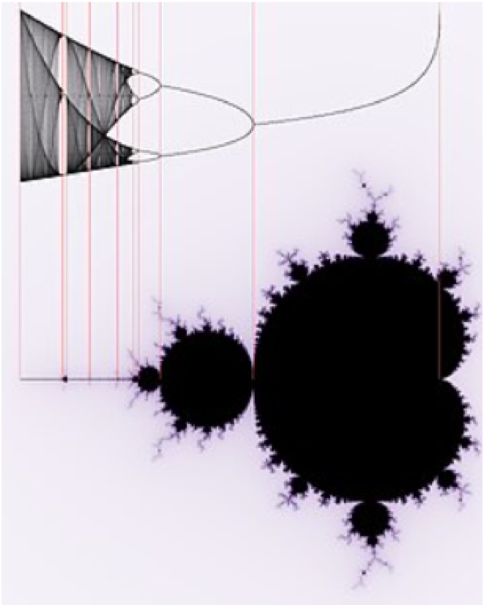

Figure 8: The Mandelbrot set, with each bulb connected to its respective bifurcation branches.

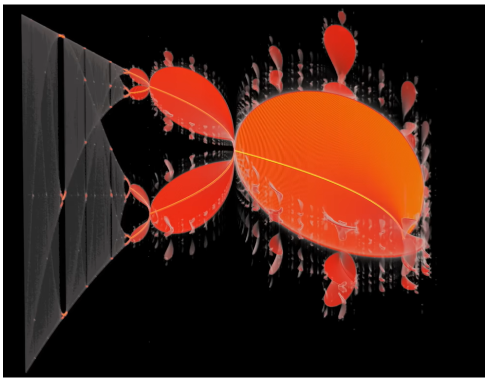

Figure 9: The Mandelbrot set bifurcating at each bulb.

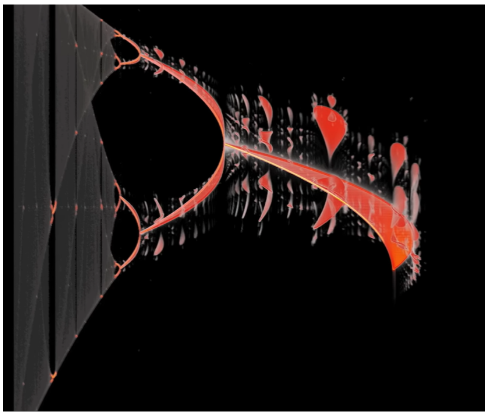

Figure 10: The Mandelbrot set bifurcating at each bulb.

Conclusion

What chaos ultimately reveals is a different kind of beauty. Not the clean symmetry of textbooks, but something messier, more alive, the kind that only appears when you look closely enough, and then refuses to stop appearing. From weather systems to ski slopes to abstract maps, the same mathematics keeps surfacing. Simple equations, behaving in ways that are anything but. This is what makes chaos unsettling and wonderful in equal measure. It doesn't emerge from complexity but from simplicity. Three coupled equations gave Lorenz an infinite, never-repeating shape. A single quadratic map gave Feigenbaum a universal constant hiding across chemistry, circuits, and populations alike. The universe keeps finding the same patterns. It isn't random. It isn't tidy. It is something in between. We're used to thinking of mathematics as the subject of certainty. But chaos directly contradicts that. Change something tiny at the start, and suddenly you're somewhere completely different. Ironically, Lorenz merely went to get a coffee and perhaps ended up changing history. The universe is neither ordered nor disordered. It is both, always, at every scale. And that, it turns out, is exactly enough.