Chaos and Intrinsic Unpredictability

Jeffrey Heninger and Aysja Johnson

31 March 2023

Contents

1 Introduction

Chaos restricts what predictions can be made by an arbitrary intelligence. These results of chaos theory are intrinsic to the dynamics itself, and are not merely statements about our current inability to understand things. While many things are not chaotic, and some chaotic things are avoidable or controllable, understanding chaos theory is important for understanding how intelligence relates to the physical world.

The idea of unpredictability itself has some subtleties. Different questions you ask about the future might be unpredictable in different ways. I discuss different features of unpredictability in Section 2: *What is Unpredictable?* Chaos theory is one of the multiple underlying reasons that something could be unpredictable. Other sources of unpredictability include computational irreducibility, anti-inductive systems, and random variables. In Section 3: *Sources of Unpredictability*, I state a definition for each one, give a few examples or counterexamples, and briefly say how they restrict predictability. This is not meant to be an exhaustive list of all of the reasons why something could be unpredictable, or a general theory of how unpredictability works. Instead, it is a list of some types of uncertainty that we know exist. Chaos theory is the best understood of these notions of unpredictability, both by humanity as a whole and especially by me in particular, so I focus this report on it.

The definition of chaos given in Section 3.1 contains multiple types of complicated motion. Section 4: *Varieties of Chaos*, contains a description of some of the important distinctions between these types of chaos. The first important distinction is between models of chaos which involve discrete time vs continuous time. Chaos can occur in either situation, and there is a natural connection between the two called a Poincaré return map. The definitions of the terms 'bifurcation' and 'multistable' are given. I then briefly describe what it means for a chaotic system to be ergodic, for it to have a strange attractor, and for it to be Hamiltonian. Finally, I define turbulence, the commonly occurring form of chaos found in fluids.

Once I have established some understanding of what chaos is and can be, I describe how to tell if something is chaotic, whether for a mathematical model or from experimental data. These techniques are described in Section 5: *Identifying Chaos*, and are straightforward to use, especially if the chaotic system can be fully described using only a few variables. If the chaos requires many variables to describe, it is often difficult to use experimental data to distinguish it from noise.

Engineers sometimes encounter something chaotic when designing something, and it seems interesting to look at how they deal with it. The answer, briefly described in Section 6: *Engineering Around Chaos*, is that they typically try to get rid of it. If possible, engineers will remove or constrain chaotic parts to make design easier. If this is impossible, the engineers' response is typically to build an analogue model and measure what it does. There are also a few things which are designed to be chaotic, such as random number generators.

I started this report talking about what sorts of predictions are and are not possible when there is chaos. Section 7: *Following Trajectories and Taking Averages* investigates what we can and cannot predict in more detail.

The most valuable prediction is usually: Given what I know about the present, what exactly will happen in the future? This is a question about a particular instance (What will happen to me?), not a statistical or counterfactual prediction. This kind of precise prediction is often impossible.

Since chaos can make predictions of a particular instance impossible, we might instead settle for statistical predictions. What can I predict will happen on average and with what shape of probability distribution? Even though this does not tell me what will happen in particular, it can provide lots of information about the future. However, there can also be problems here: the average might change chaotically and be unpredictable itself. In order to make reliable predictions, averaging should be done over a scale which is larger (in space or time) than the chaos. People have developed many ways to make these sorts of statistical predictions, some of which are much more mathematically sophisticated than others: experimental analogues, direct numerical simulations, current machine learning techniques, leveraging invariant quantities, and periodic orbit theory, all of which will be explained in Section 7.

Even when you cannot predict the motion of a particular trajectory or the statistics of trajectories, there is often some predictable pattern in the chaos. Most of the information about the future is unpredictable, but not quite everything. The predictions you can make are often far removed from the original question you were trying to answer. A few examples will illustrate this and are described in Section 9: *Other Predictions*. Turbulent spectra do not tell you what particular motion of a fluid will occur in the future, nor the statistics of this motion, but they will tell you the relationship between the statistics at different scales: small waves or whirls in the fluid's motion systematically contain less energy than larger ones. Feigenbaum constants do not tell you where a transition to chaos occurs, but they will tell you that for this particular route to chaos, there is a kind of self-similarity. There are lots of interesting patterns to discover within chaos, even when predicting individual trajectories is impossible.

People have also worked on ways to control chaos. The chaos seems to be able to do a lot of different things on its own. Can you pick which of those things actually happen, using some inputs to the system? The answer is sometimes yes, although restrictions do apply. How this can be done is described in Section 8: *Controlling Chaos*.

This report serves as reference material for several pages on the AI Impacts Wiki,[1, 2] as well for several blog posts.[3, 4, 5] The goal of the overall project is to explore the relationship between chaos and intelligence, both for humans and for future artificial intelligence. While this project is the reason why I wrote this report, it should be a helpful introduction to chaos theory, regardless of what motivated your interest.

2 What Is Unpredictable?

Unpredictability is not typically a binary, cut and dry property of a system. The predictability of a system does not even always vary monotonically. There are multiple aspects of a system that one might want to predict and multiple things that can influence whether the answers are predictable.

Firstly, the magnitude of the uncertainty can vary, with different questions being answerable with different levels of precision. If the best prediction is less precise, or has a probability distribution with a larger range, then one can be more wrong about their prediction. Moreover, sometimes it is possible to reliably predict the probability distribution, but sometimes the distribution itself can vary unpredictably.

Secondly, how predictable some aspect of the system is might depend on the time and scale you view the system with.

Prediction ability tends to be worse the farther out one is predicting. For instance, estimating who will be elected president in twenty years is much more difficult than estimating who will be elected next month. This decline in predictability can be attributed to a few causes. For one, the further out in time you go, the more noise will tend to swamp the prediction. Another is that enumerating all possible trajectories is computationally expensive and the branching increases exponentially as time progresses. In light of these difficulties, it seems natural to assume that predictive ability will strictly decline over time, i.e., that it will monotonically decrease.

Strangely, though, this is not always the case. For instance, one can predict the weather on relatively short time horizons (it will rain tomorrow) *and* on some very long time horizons (ice ages), but intermediate predictions (a few weeks out) are difficult. What scales something is or is not predictable on depends on the details of the thing you are studying.

The scale at which one views the system also matters. Predicting the movement of atoms is different from predicting the changes of macroscopic quantities. Different physical systems are involved, which may or may not be predictable. As you zoom out from small scales to larger scales, the system might get less or more chaotic and so its predictability will change. Predicting the motion of a fluid at the atomic scale is hard, predicting the motion of a fluid at the micrometer scale is easy, and predicting the motion of a fluid flowing rapidly through a pipe is hard.[1] Coarse-graining a system sometimes makes it more predictable and sometimes does not.

[1] I talk about this more in Section 7.5.

3 Sources of Unpredictability

There are different reasons that a system might be unpredictable. The first is that it could be *chaotic*, i.e., any uncertainty in the initial conditions grows exponentially until it reaches maximum uncertainty. It could be *computationally irreducible*, meaning that there is no way to figure out the output of a computation faster than running the computation itself. It could be *anti-inductive*, meaning that there are incentives to keep a system unpredictable. Finally, it could be modeled as a *random variable*, which allows us to make statistical predictions even when we do not know what in particular will happen.

3.1 Chaos

Definition

Chaos theory describes how a deterministic system (typically a differential equation or a difference equation) can be unpredictable, because any initial uncertainty grows exponentially in time. Edward Lorenz, one of the fathers of the field, described chaos as:

> When the present determines the future, but the approximate present does not approximately determine the future.[6]

There is no universally accepted definition of chaos. Most definitions involve these two characteristics:

1. **Sensitive Dependence on Initial Conditions.** In a chaotic system, if you start with two initial conditions which are extremely close together, the distance between the trajectories grows exponentially in time. Equivalently, if you start with a small amount of uncertainty around an initial condition, that uncertainty grows exponentially in time. The growth rate is called the *Lyapunov exponent*.
2. **Recurrence / Mixing / Topological Transitivity.** The details of what exactly is required varies for the different definitions, but that should be some way for the motion to return to places close to where it has been before. The motion does not just spread out: it returns and experiences the sensitive dependence on initial conditions again and again.

An intuitive way of thinking about these two conditions is 'Stretch and Fold'. The sensitive dependence on initial conditions stretches nearby trajectories apart from each other, while folding allows the trajectories to come back close to where they have been before.

An alternative definition of chaos is that the number of qualitatively different trajectories grows exponentially with time. This growth rate is called the *topological entropy*. Something which has a positive topological entropy will also satisfy the other two conditions for chaos. This definition is harder to make precise mathematically, so I will not use it here.

Counterexamples

Intuitively, chaos refers to motion that is complicated. Both parts of the definition are needed so that certain kinds of simple motion are not considered chaotic. Here are some counterexamples:

- **Pencil Balanced on its Tip.** Small changes to the initial position of the pencil determine which way it will fall. Once it falls, it just sits on the table. This has sensitive dependence on initial conditions, but it is not recurrent.

  The differential equation $\frac{dx}{dt} = x$ is similar. All initial conditions move apart exponentially, but there is no recurrence.
- **Periodic Motion.** Something which oscillates periodically is recurrent, but nearby trajectories do not separate. It does not have sensitive dependence on initial conditions.
- **Brownian Motion.** The motion of a particle in a fluid, subject to collisions with the molecules of the fluid, is recurrent and the uncertainty in its position grows, but it is not chaotic. The uncertainty grows as the square root of time, $\Delta x \sim t^{1/2}$, which is much slower than exponential.

Examples

Here are a few examples of chaotic systems to help you better understand the definition. More examples will be discussed throughout this report.

- **Shift Map.** Write any number between 0 and 1 as a decimal. At each time step, take all of the digits and shift them left one place. Drop any number to the left of the decimal to stay between 0 and 1. This can be expressed as the difference equation: $x_{n+1} = 10x_n \mod 1$. Rational numbers are periodic, while irrational numbers are aperiodic. Two similar numbers grow 10 times farther apart with each time step, so this has sensitive dependence on initial conditions. The mod keeps everything between 0 and 1, so this is recurrent. This is chaotic and almost all initial conditions move aperiodically.
- **Kneading Dough.** This is the most literal example of stretching and folding. Two dots of food coloring in the dough get farther apart when you roll out the dough. Folding the dough back onto itself makes it possible for a dot of food coloring to return to close to where it came from. The patterns of marbled bread are reminiscent of the patterns you see in chaos theory.
- **Pinball.** Consider a pinball bouncing between the three disks near the top of a pinball table. The curvature of the disks means that the trajectories will diverge as the result of these collisions. The pinball can return close to where it started, so it is recurrent.[2]

Limitations on Prediction

In a chaotic system, any initial uncertainty grows exponentially: $\Delta x \sim \Delta x_0 \, e^{\lambda t}$, where λ is the Lyapunov exponent. Rearrange this to determine how long you can make predictions for:

$$t \lesssim \frac{1}{\lambda} \log \left(\frac{\Delta x_T}{\Delta x_0} \right).$$

How long you can make predictions only depends on a few quantities:

- **Lyapunov Exponent**, λ. This is a property of the system you are studying. Chaotic systems have a positive Lyapunov exponent, while non-chaotic systems have a Lyapunov exponent of zero.
- **Initial Uncertainty**, Δx_0. If you can measure your initial conditions more precisely, you can make predictions for a longer amount of time. The initial uncertainty can come from measurement error, manufacturing error, theoretical approximations, numerical resolution, outside influences that are not accounted for, or thermal noise. Even if all classical uncertainties can be reduced to zero, there is still the fundamental uncertainty of quantum mechanics, which imposes a hard limit beyond which prediction is impossible.
- **Tolerable Uncertainty**, Δx_T. How precise of a prediction are you trying to make? It is harder to make more precise predictions, so you can make less precise predictions for longer.

Unless the initial uncertainty and tolerable uncertainty are orders of magnitude apart, this is approximately $t \lesssim 1/\lambda$. $1/\lambda$ is called the Lyapunov time.

It is possible to predict a chaotic system up to the Lyapunov time. Making predictions longer than this gets exponentially harder, and then you run into the limits imposed by quantum mechanics. For longer times than this, it is sometimes possible to make some statistical predictions for what will happen, but not to predict which particular trajectory will be followed.

3.2 Computational Irreducibility

Definition

In computationally irreducible systems, the output of a computation cannot be obtained any faster than simply running the computation itself. In other words, computationally irreducible systems are incompressible.

[2] I wrote a detailed calculation of the Lyapunov exponent for a game of pinball in You Can't Predict a Game of Pinball.

This idea is due to Stephen Wolfram, who describes it as:

> Whenever computational irreducibility exists in a system it means that in effect there can be no way to predict how the system will behave except by going through almost as many steps of computation as the evolution of the system itself.[7, 8]

Wolfram goes on to argue that all computationally irreducible processes perform at the maximum level of computational power and so are computationally equivalent[9] to each other. While I am skeptical of the Principle of Computational Equivalence, I find the notion of Computational Irreducibility to be an interesting potential source of unpredictability.

Computational irreducibility is distinct from chaos. In chaotic systems, one could potentially predict the outcome, but only if one had infinite precision over the initial conditions. Computational irreducibility, on the other hand, only requires a finite amount of precision, because it occurs in a discrete system. It is possible to re-initialize or exactly copy the current state of the system, and the behavior will be exactly the same as the original run. Chaotic systems cannot be returned to the exact same starting conditions or exactly copied using a finite amount of precision.

Examples

- **Cellular Automata.** Much of Wolfram's motivation for computational irreducibility seems to come from investigations of cellular automata, especially Class IV, which involve a mixture of order and disorder. Locally, the structure is often fairly simple, but the way the simpler substructures move around and interact is complicated. The details of these movements do not seem to be predictable by any simpler program, although some coarse-grained information can be predicted.[10]
- **Pseudorandom Number Generators.** These are algorithms used to create sequences of numbers with similar properties to sequences of random numbers. The algorithm is not truly random, but is instead calculated from an initial value called the seed. This is re-initializable: given the same initial seed, the algorithm will produce the same sequence of numbers. Pseudorandom number generators are used extensively for Monte Carlo simulations, procedural generation of computer graphics, and cryptography.

Limitations on Prediction

If there is computational irreducibility, prediction is possible, but only by performing an equivalent calculation. It is impossible to compress the computation so you can make predictions using less cost than the computation itself requires. It is possible to exactly learn the future before it happens, if you have something that can perform the computation faster. You might also want to make a prediction if the two computations are distinct for other reasons: for example, a perfect flight simulator would be useful[3] because crashing a simulated airplane is less bad than crashing a real airplane.

3.3 Anti-Inductive Systems

Definition

Something is anti-inductive if the act of noticing and responding to a pattern also destroys that pattern. "Whatever is believed in, stops being real."[11] Regularities are disincentivized in anti-inductive systems, either because the actors within such systems have reason to be unpredictable, or because regularities are quickly exploited until they no longer exist.

Anti-inductive systems appear when there are many intelligent beings competing with each other in a zero sum game, with a unified notion of success for everyone. Each individual attempts to identify patterns in everyone else's behavior that they can exploit to gain an advantage. Once a pattern is observed and becomes widely known, the individuals adjust their behavior in a way that destroys this pattern.

[3] If it is possible to create.

If there are too few players, then the best strategies might not be used enough to destroy the pattern. If there are many different notions of success, then people aren't competing directly with each other and so are less likely to exploit each others' strategies.

Many forms of human interaction are a mix between being anti-inductive and not, or are anti-inductive in some circumstances but not in others. I will address examples and counterexamples together.

Examples and Counterexamples

- **Stock Market.** The classic example of an anti-inductive system is the stock market. Lots of people are trying to use the stock market to make money, so we should expect that any observed patterns in the stock market will immediately vanish. If stock prices are systematically lower on Friday and higher on Monday, then people are incentivized to buy more stocks on Friday and sell more stocks on Monday. This raises Friday's stock prices and lowers Monday's stock prices, removing the pattern.[12] The stock market is not a zero sum game: the market as a whole tends to increase. It is possible to make money on the stock market by not taking part in the competitive game and instead investing in a mutual fund. In order to reliably perform better than the market, you need faster trading, insider information, or knowledge of a pattern that other people have not yet noticed. This is the efficient market hypothesis.
- **Poker (or other deception games).** One of the most important skills in poker is being able to estimate an opponent's hand. Thus, being unpredictable in one's own plays is critical for success.
- **Job Applications.** Job applications have also been mentioned as an anti-inductive system.[13] Applicants want to stand out, so they are more likely to be selected. Trying to stand out from a group of people is anti-inductive: when people learn what has worked in the past, they copy it and it ceases to make them stand out. Some job applications seem more anti-inductive than others. If the application is more competitive and standardized, then standing out becomes more valuable and the process is more anti-inductive.
- **Dating.** Dating also involves a potentially large number of people interacting in a potentially competitive situation. Dating apps, in particular, bring lots of people together in a single system, which makes the competition more explicit and makes it more important to stand out. Dating in person is much less anti-inductive. There is a much smaller group of people to compete with and success is more personal.

Limitations on Prediction

In order to perform better than average in an anti-inductive system, it is not enough to make good predictions or to design good strategies for yourself. You have to make better predictions or strategies than anyone else. Patterns have existed in the stock market, and people have made money taking advantage of them, but only as long as they did not become common knowledge.

Assuming that some undiscovered patterns still exist in the stock market, a superintelligence could recognize them and make money off of the human traders. However, if there were multiple superintelligences trading in the stock market, with sufficiently large liquidity, competing against each other, this would not work. The interactions between the superintelligences would render the behavior of stock prices inscrutable even to them.

3.4 Random Variables

Definition

A random variable is a model of something for when you do not know its behavior. The random variable allows you to make statistical predictions, even when you cannot make a prediction about what in particular will happen in the future.

Modeling something as a random variable is a tool for dealing with uncertainty. It does not explain why the thing you are studying is uncertain. Two things modeled by the same random variable might be uncertain for different reasons. I will now describe several possible causes of randomness, followed by some examples, and then the limitations on prediction.

Causes

- **Intrinsic Randomness.** Something might be random because the relevant laws of physics are themselves random. In this case, the behavior of the system is not uniquely determined by its past and the best possible predictions are statistical.
- **Amplified Randomness.** Macroscopic objects are typically not directly affected by intrinsic randomness. However, if they have sensitive dependence on initial conditions,[4] they can quickly amplify any microscopic uncertainty to the macroscopic scale.
- **Epistemic Randomness.** Random variables can also be used to reflect your lack of knowledge, instead of being a claim about randomness of reality itself.

Showing that something can be modeled well using a random variable does not tell you what mechanism gave rise to the randomness. One possibility is that microscopic uncertainty was amplified by chaos, although other explanations are possible. It is best to understand randomness and chaos as distinct, even though they have similar colloquial meaning and share some history.[5]

Examples

- **Dice & Galton Boards.** Many physical random number generators work by amplifying smaller uncertainties. In a Galton board, which side of each peg the ball bounces towards depends sensitively on the ball's initial position. The bouncing of dice off of their edges similarly amplifies uncertainty, especially if the surface they are rolling on is rough.
- **A Shuffled Deck of Cards.** This is epistemic randomness. It is best modeled as a random variable because you do not know what order the cards are in. The process of shuffling the cards mixes them and, because you do not perfectly alternate cards from the two stacks, introduces uncertainty. Once the deck of cards has been shuffled, the uncertainty is entirely in our knowledge about them.
- **Quantum Measurement.** Solving the equations of motion for a quantum system results in a complex wave function. To connect this wavefunction with experimental data, we take its absolute value squared, which gives the probability of finding the particle in each particular state.
- **'Molecular Chaos.'** The distinction between randomness and chaos is most blurred with the term 'molecular chaos,' which predates the development of chaos theory.[6] Molecular chaos refers to the hypothesis that the motions of particles in a gas are independent of each other. If we model the particles as hard spheres, then this is a form of billiards with lots of collisions involving curved surfaces, so this is chaotic. Two initial conditions which are closer together than the diameter of these hard spheres will diverge exponentially. However, we look at the motion at a larger scale, so the particles are points and each collision results in a random change in the velocity. This is Brownian motion and it is not chaotic.

Limitations on Prediction

Random variables allow you to make reliable statistical predictions, even when you cannot make a reliable prediction of a particular instance.

[4] Recurrence is less important here.

[5] Prior to the development of chaos theory in the mid-1900s, the word 'chaos' was used to refer to completely disorganized matter.

[6] It seems as though the term was coined by Huxley in 1869 to refer to something like the heat death of the solar system.[14] He was likely influenced by Maxwell's recent work on the kinetic theory of gases.[15]

Showing that something can be well modeled using a particular random variable does not guarantee that the particular instances are unpredictable. To determine if better predictions are possible, you have to determine what caused the randomness. If the randomness is epistemic, then it could be possible to learn more about it. If the randomness is intrinsic, then no better predictions are possible. If the randomness is because of amplified uncertainty, then the restrictions from chaos theory apply. If the randomness is caused by something else, then you need to determine what that something else is before claiming to have found a limitation on prediction.

3.5 Summary

There are multiple different ways that and reasons why something can be unpredictable. It is worth remembering these other forms of unpredictability, even though the rest of this report will be focusing on chaos theory.

The main difference between chaos and computational irreducibility is that computational irreducibility is completely discrete, while chaos has continuous space and/or time.[7] The initial conditions for something that is computationally irreducible involve a finite amount of precision, while the initial conditions for something that is chaotic involve much more precision than is possible. This makes it possible to re-initialize or copy a computationally irreducible system, and its subsequent behavior will be exactly the same. Chaotic systems are impossible to copy or re-initialize.

For a system to be anti-inductive, it has to involve agents which are looking for patterns. When a pattern is first discovered, agents can profit off of the pattern. However, the act of exploiting the pattern also destroys the pattern. Anti-inductive systems might be chaotic, or have sub-components which are chaotic, but it is difficult to show this. If you discovered any evidence that the system followed some equations of motion that are chaotic (for example), someone would use this knowledge to profit off of the system and destroy the pattern.

Random variables are a tool for dealing with uncertainty. Sometimes, physics is intrinsically random, and it is impossible to make reliable predictions. Random variables can also be used to describe systems which are not random themselves, but which you are uncertain about. Showing that something is well described by a random variable does not tell you if it is possible to make better predictions unless you can identify the source of the randomness.

4 Varieties of Chaos

4.1 Discrete vs Continuous Time

Time is central to the definition of chaos. Both the Lyapunov exponent and the topological entropy are defined as the rate at which some quantity grows exponentially over time. It is not surprising that discrete and continuous time result in different kinds of chaos.

Most of the most successful laws of physics take the form of differential equations: equations that relate the time derivative of some quantity to the quantity itself or some other quantities. Each quantity whose time derivative is involved in the differential equations is called a 'dynamical variable.' The space of all possible combinations of the dynamical variables is called the 'state space.' Time varies continuously and can take the value of any real number. A trajectory is a continuous curve through state space.

Because of continuity arguments, a differential equation with only one or two dynamical variables cannot exhibit chaos: you need at least three variables for the trajectories to loop around past each other in complicated ways. One of the simplest examples of chaos in continuous time is the Rössler Equations:

[7] At least to the resolution we have available.

$$\begin{aligned}\frac{dx}{dt} &= -y - z \\ \frac{dy}{dt} &= x + ay \\ \frac{dz}{dt} &= b + x(z - c)\end{aligned}$$

A typical trajectory moving through three-dimensional space is shown in Figure 1.

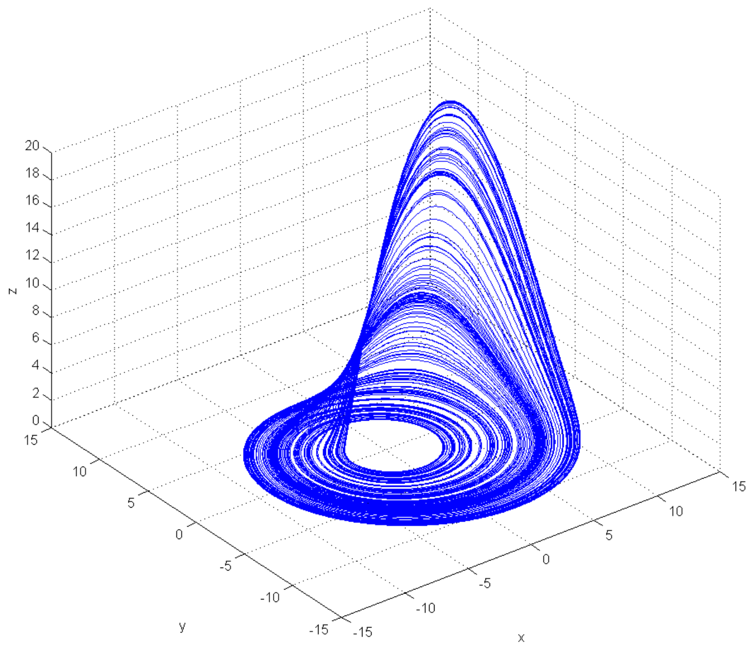

Figure 1: A typical trajectory for the Rössler Equations with $a = b = 0.2$, $c = 5.7$. One thing I particularly like about this system is that it is easy to see how, as trajectories go around counterclockwise, the ones near the edge are stretched upwards and then folded back down. *Made by Jeffrey Heninger.*

For a discrete time system, time is an integer, not a real number. The value of some quantity at the next time step, x_{n+1}, is a function of that quantity or some other quantities at the previous time step. These are called 'difference equations,' 'recurrence relations,' or 'maps.' A trajectory is a discrete set of points in state space.

A difference equation only needs to have one dynamical variable to be chaotic. The continuity arguments for differential equations do not apply because the orbit can jump from point to point. One simple and historically important example is the logistic map:

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

Its motion on the real number line between 0 and 1 is shown in Figure 2.

There is a natural way to connect continuous and discrete motion, called a Poincaré return map.[8] In a continuous time system, choose a surface which is transverse (nowhere tangent) to the motion. Follow the

[8] The surface is called a Poincaré section. See also Wikipedia: Poincaré map.

Start

Figure 2: Two trajectories of the logistic map with $r = 4$ which begin close together. *Made by Jeffrey Heninger. If the gif does not work, try a different pdf viewer.*

motion forward in time, but only record the locations where the motion crosses the surface. Recurrence guarantees that the motion will continue returning to this surface. The result is a discrete time chaotic system, the Poincaré return map, which corresponds to the original continuous time chaotic system. Figure 3 shows how this works for the Rössler attractor.

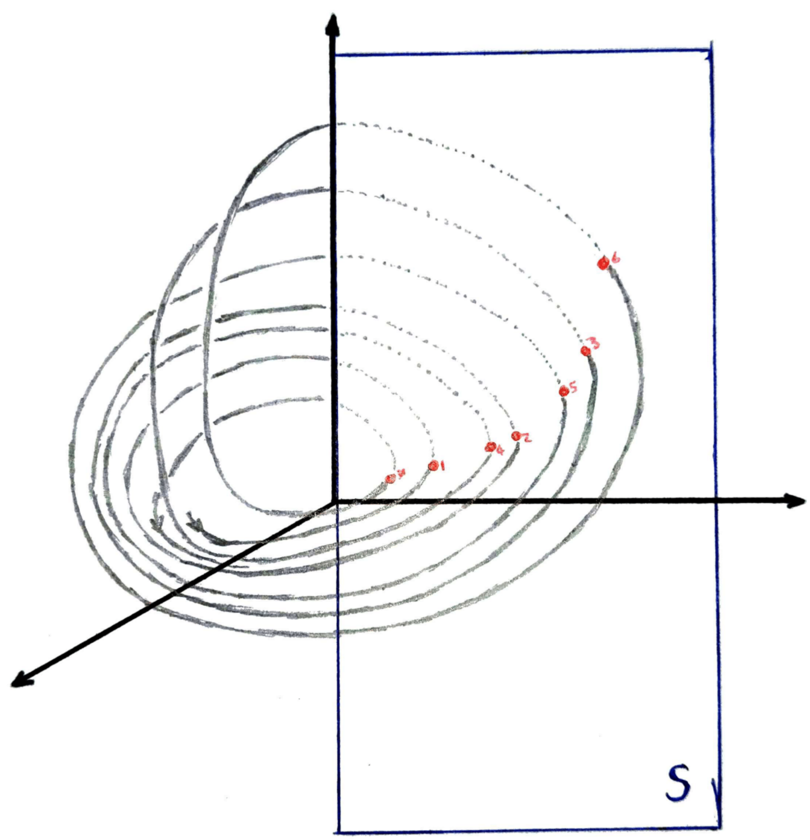

Figure 3: (black) Coordinate axes of the state space for the Rössler Equations. (blue) A Poincaré section for the Rössler Equations. (gray) A trajectory in the original state space. (red) A trajectory of the Poincaré return map. *Drawn by Jeffrey Heninger.*

4.2 Bifurcations and Multistability

Dynamical systems often depend on some parameters that characterize the system. The motion of the pinball depends on the radius of the disks and the distance between them. The motion of a double pendulum depends on the lengths and masses of each pendulum. The flow of water through a pipe depends on the pressure difference between where the water enters and leaves the pipe.

You might expect that small changes in the parameters results in only small changes to the motion. This is almost always true, for non-chaotic systems. However, there often are a few parameter values where the resulting motion changes discontinuously, and qualitatively different motion occurs. Water flows smoothly through a pipe as long as the pressure difference is small enough, while when the pressure difference is larger, there are puffs of turbulence. An old car might only rattle if you drive faster than a certain speed. The transition to qualitatively different motion at a parameter value is called a bifurcation.

The study of bifurcations is an important part of dynamics. There are only a few types that occur commonly, like saddle-node bifurcations, Hopf bifurcations, and period doubling bifurcations. Understanding bifurcations helps us understand how things can transition to chaos. The details are not important to us here, but can be found in many textbooks on dynamics.[9]

If qualitatively different behavior occurs depending on how you start the motion, this is not a bifurcation. Instead, it is an instance of multistability (or bistability). The long time behavior of the motion is different for different initial conditions. The set of all initial conditions which approach a particular long-time behavior is called a basin of attraction. The boundaries of basins of attraction can be extremely complicated and an initial condition near them often behaves chaotically for a little while before settling down into one of the basins of attraction.

4.3 Ergodicity

An ergodic system is something for which almost all trajectories eventually explore all parts of the space they are moving in. This term can be used for either chaotic motion or stochastic motion. The chaos (or randomness) extends to all parts of the space.

The shift map is one example of ergodic chaos. Almost all real numbers between 0 and 1 eventually have every digit appear in their decimal expansion (and every pair of digits and . . .). As the digits shift left, these points will eventually be in every part of the interval.

Many variations of billiards with curved walls, including Sinai billiards and the Bunimovich stadium, shown in Figure 4, are ergodic. The motion of molecules in an ideal gas is ergodic as well.

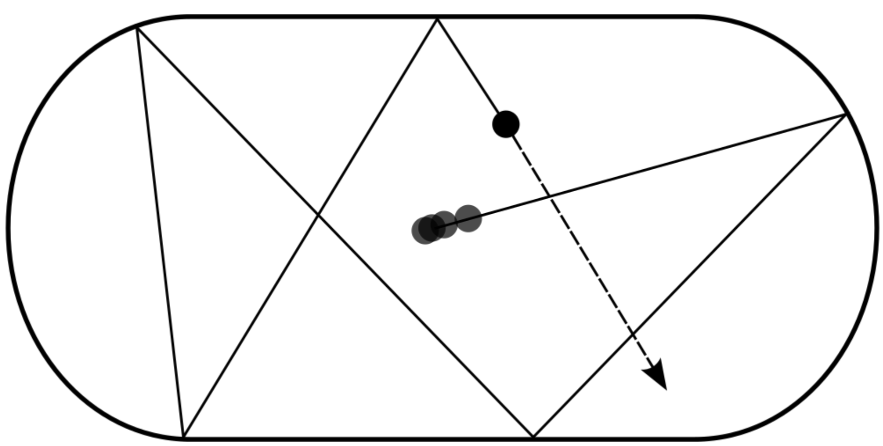

Figure 4: The motion of a billiard ball in a Bunimovich stadium. *Source: Wikipedia Commons.*

[9] For example, Chapters 3 & 8 in Strogatz.[16]

4.4 Strange Attractors

Chaotic systems with a strange attractor are not ergodic. Instead, almost all initial conditions will approach a subset of the space, called the attractor, and not return to near their starting location.[10] An attractor is strange if the motion on it is not stationary or periodic.

Strange attractors typically have a fractal shape, because of the repeated stretching and folding of the dynamics. Each time the dynamics stretches and folds, it creates another fold in the attractor. After a long time, the result will be a shape that has lots of folds nested inside each other, much like the Koch snowflake or dragon curve.

Strange attractors look cool, so they are the most commonly seen pictures of chaos. I have already shown the strange attractor for the Rössler Equations. Figures 5 and 6 show two more strange attractors, one with discrete time and one with continuous time.

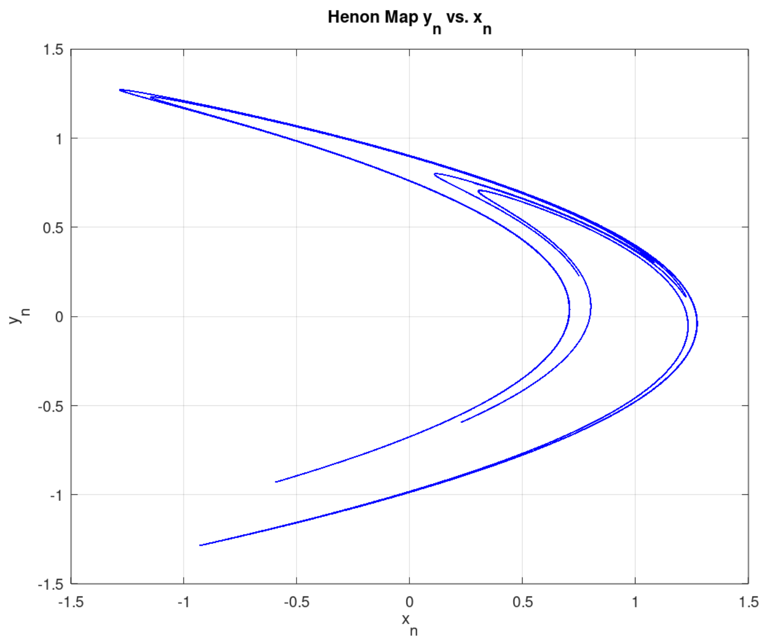

Figure 5: The strange attractor of the discrete time Hénon Equations. *Made by Jeffrey Heninger.*

4.5 Hamiltonian Chaos

Strange attractors typically occur when energy is being lost to friction and when additional energy is somehow being added. If energy is being conserved instead, the system is usually Hamiltonian. Hamiltonian systems conserve state space[11] area, so it is impossible for lots of different initial conditions to converge onto an attractor.[12]

The chaos that occurs in Hamiltonian systems is very constrained and looks quite different from other kinds of chaos. It typically has 'island chains' separated by bands of chaos. Hamiltonian chaos is often not ergodic because the chaotic trajectories cannot enter the islands.

[10] One of the reasons why people use different definitions for recurrent / mixing / topologically transitive is because people typically want something which is recurrent on the attractor, but not in the whole space, to count as chaotic.

[11] For Hamiltonian systems, state space is typically referred to as phase space. It usually involves an equal number of position and momentum variables.

[12] In order to converge onto an attractor, a bunch of different trajectories have to be getting closer together. If you draw a circle around them and let the circle move with the trajectories, then the circle will shrink. This contradicts conservation of phase space area.

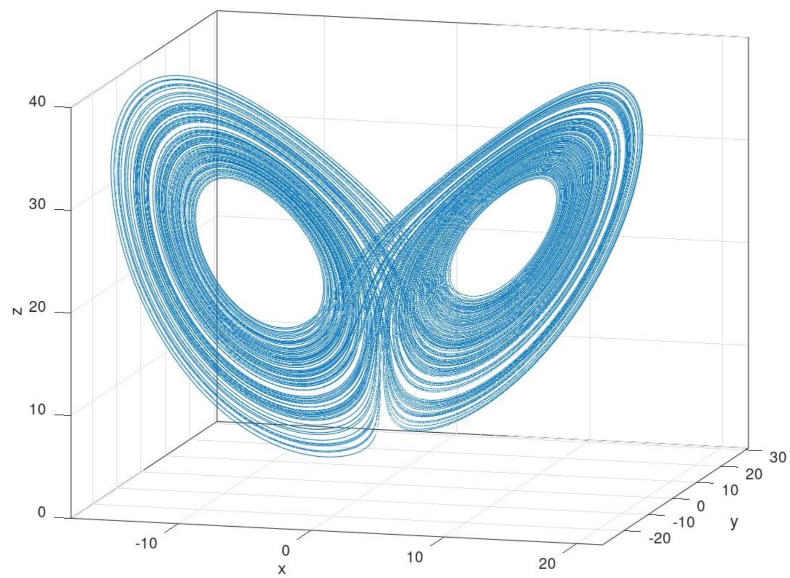

Figure 6: The strange attractor of the continuous time Lorenz Equations. *Made by Jeffrey Heninger.*

One of the most thoroughly studied[13] examples of Hamiltonian chaos is the Standard Map, which is a model of a pendulum which is periodically kicked (or hit with a hammer) that has been used to describe many physical systems:[14]

$$\begin{aligned} p_{n+1} &= p_n + k \sin\theta_n \\ \theta_{n+1} &= \theta_n + p_n + k \sin\theta_n \end{aligned}$$

Figure 7 shows trajectories of the Standard Map, which help to show how motion in the phase space of a Hamiltonian system is qualitatively different from motion which approaches a strange attractor.

4.6 Turbulence

So far, the examples of chaos I have given involve a few quantities which change as a function of time. Turbulence involves fields,[15] which is something that can be measured at any location and can have a different value at different places. For example, the pressure, temperature, and velocity of a fluid are fields – and can be turbulent.

Turbulence has even more definitions than chaos. One useful definition is that turbulence is something which is chaotic in both time and space. Two instances of turbulence which are similar at one location will become more distinct exponentially as you move to other locations. Patterns which appear at one location can be similar to independent patterns which appear at other locations. Another much less precise, but still useful definition of turbulence is: Something like a fluid looks like it is doing something really complicated. Figure 8 shows a drawing by Leonardo da Vinci of turbulence in water.

[13]The most thoroughly studied example of Hamiltonian chaos is probably the Three Body Problem of celestial mechanics.

[14]Things modeled by the Standard Map include: confinement of charged particles in a magnetic mirror,[17] motion of charged particles in a particle accelerator,[18] orbits of comets,[19], and ionization of Rydberg atoms.[20]

[15]Wikipedia has a decent introduction to what fields are in physics: Field (physics).

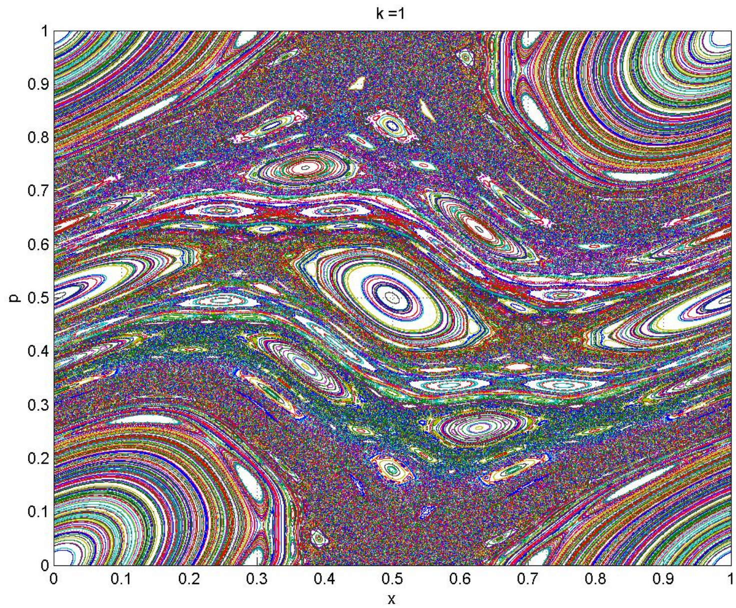

Figure 7: The phase space of the standard map with $k = 1$. Since trajectories starting at different locations have qualitatively different behavior, many trajectories are shown. Each trajectory is a different color. *Made by Jeffrey Heninger.*

Figure 8: Water flowing from a sluice into a pool by Leonardo da Vinci.[21]

5 Identifying Chaos

There are two situations in which you might want to tell if something is chaotic: (1) you have a mathematical model and you want to see if it exhibits chaos, or (2) you have some experimental data, usually a time series, and you want to see if the thing that produced it is chaotic. For both, it is usually easy to tell if the motion is recurrent and the challenge is determining if there is a positive Lyapunov exponent.[16]

Rigorously proving that a mathematical model is chaotic is difficult, but often possible. Sharkovskii's Theorem proves that if a one-dimensional discrete time map has a periodic orbit of period three, it must have periodic orbits of every other period, and arbitrarily complex motion is possible.[22, 23][17] Lorenz was able to prove that his equations do not approach any periodic motion after an arbitrarily long time.[24] The KAM Theorem describes how periodic motion can become chaotic in Hamiltonian systems.[25, 26, 27][18]

Numerically estimating the value of the Lyapunov exponent and showing that your estimate is positive is easier than rigorously proving that there is chaos. Take two initial conditions extremely close together and numerically integrate the equations of motion to find approximate trajectories. If the two trajectories move apart exponentially, then this is numerical evidence that the mathematical model exhibits chaos.[19]

Lyapunov exponents can also be estimated from experimental time series.

The first step is to reconstruct the state space using delay coordinates. Instead of just plotting the time series itself, $x(t_i)$, treat neighboring time steps as independent variables and plot $x(t_i)$ vs $x(t_{i-1})$. For example, if the time series is of position, you might expect that the state space involves position and velocity. The time delay coordinate is a second variable which combines position and velocity: $x(t_{i-1}) \approx x(t_i) - \Delta t\, v(t_i)$. The space spanned by $x(t_i)$ and $x(t_{i-1})$ is the same as the space spanned by $x(t_i)$ and $v(t_i)$. If introducing a second variable is not enough, you add additional neighboring time steps to produce a larger reconstructed state space. You can continue adding time-delayed variables until any two points in your time series which are close to each other in the state space initially move in about the same direction.

Once you have the reconstructed state space, look for an instance where the time series returns close to a location in the reconstructed state space. Treat the original point and the recurrent point as distinct initial conditions. How does the distance between the two trajectories subsequently change? If the distance grows exponentially, then there is a positive Lyapunov exponent and the experiment was chaotic. The details of how this works, along with a good description of Lyapunov exponents and why we care about them, can be found in this paper by Wolf et al.[28]

This procedure works well when the chaos is low-dimensional: when the minimum number of variables needed to describe the chaos is small. If the chaos is high-dimensional, then there is often not enough data, or not precise enough data, to reconstruct the state space. These experiments are often indistinguishable from noise. When we are confident that something exhibits high-dimensional chaos, it is often because we know what the relevant equations of motion are, like in fluid mechanics.

Once we have identified that something is chaotic, we have to figure out what to do with it.

6 Engineering Around Chaos

Most human-designed things are not chaotic.

The act of designing something involves making predictions about how it will behave. This intrinsically biases designers towards motion which is not chaotic. If there is something which is moving unpredictably,

[16] Recall from Section 3.1 that the Lyapunov exponent is the exponential growth rate of the distance between two nearby trajectories, or equivalently the exponential growth rate of the uncertainty of your initial conditions. If this does not grow exponentially, then the Lyapunov exponent is zero and the motion is not chaotic.

[17] Sharkovskii proved that there exist periodic orbits of all lengths. Li and Yorke proved this result independently, and that arbitrarily complicated motion is possible. Their paper was the first to use the term 'chaos' in the modern sense.

[18] This theorem was proved in three parts.

[19] The method described here only calculates the largest Lyapunov exponent. If you want to know whether the dynamics expands or contracts in other directions, more sophisticated numerical methods must be used.

the most common engineering's response seems to be to bolt it down or to force it to run along a track, which removes the chaos.

This seems like it is part of a larger tendency towards reducing complexity in order to make things easier to design.

There are situations where the chaos is unavoidable, because it is too difficult to prevent, too strong to ignore, or intrinsic to what needs to be done. This is most common when fluids are involved. Fluids move in many ways, and it is difficult to completely constrain their motion. The forces that can be exerted by the atmosphere or the ocean can be large enough to force us to pay attention to them. In these situations, it seems as though the engineering strategy often involves (1) making an analogue model and putting it in a wind tunnel to see what happens, (2) dampening any motion that does enter your device as much as possible, or (3) making your device tough enough to survive, despite the chaos. Engineers tend to try to reduce the effects of the chaos to something small enough that you can deal with it without having to predict exactly what will happen.

In order to make a random number generator, you actually want chaos. This is the case both for computer-generated pseudorandom number generators and for mechanical games of chance. Dice rolling and bounces off of their edges have sensitive dependence on initial conditions, especially if the surface they are colliding with is not smooth. The uncertainty of the location of a ball on a Galton board similarly increases every time it collides with a peg. For mechanical random number generators, sensitive dependence on initial conditions is much more important than recurrence, so these often do not satisfy the complete definition of chaos. For truly random numbers, people often look to nature for things where chaos or quantum mechanics provide the randomness. One fun example is Lavarand, a random number generator based on the chaotic motion of lava lamps.[29]

While human-designed things which are not explicitly trying to be random are almost never chaotic, this does not mean that human-built things are almost never chaotic. When something is built without a design, through the interactions of many people, it tends to be more complex and can often be chaotic. For example, traffic flows can have some interesting dynamics,[20] which traffic engineers try to smooth out.[30, 31, 32] The tendency to avoid chaos is because of design, not because it involves humans.

7 Following Trajectories and Taking Averages

There are two main kinds of predictions that people try to make for chaotic systems: predicting a particular trajectory and predicting the statistics of many trajectories.

The first is a prediction about the future that I will experience. I currently have a measurement of the current state of the system (the initial conditions) and some understanding of the patterns describing how things change (the equations of motion). I would like to use this knowledge to predict what, in particular, will happen in the future. Using the initial conditions and the equations of motion, it seems like I should be able to trace out what will happen (the trajectory). But to do this for a long time requires perfect knowledge of the initial conditions and equations of motion, and any uncertainty grows exponentially until these predictions are useless.

The second involves predicting the statistics of what is likely to happen. Instead of asking what will happen for these particular initial conditions, I ask what sorts of futures are most likely as a result of these equations of motion. This is often a more well-posed question to answer.

When calculating these averages, it is important to average over all of the chaos. If you instead average over only part of the chaos, then the average itself can change chaotically. This sort of chaos is often called 'non-stationary.' Predicting the future behavior of the statistical properties of the motion might also be impossible after a long enough time. It is still possible to do the procedures that calculate the averages of the motion, but the result is much less useful.

[20] The easiest of these to see are standing waves (which are not chaotic): stop-and-go patterns that often form on highways between major interchanges. Adding a few features to the model can lead to chaos.

Another way of thinking about this is that the average must occur at a larger scale (in space or time) than the chaos. Mathematical models exhibiting chaos have structure at arbitrarily small scales because the repeated stretching and folding of the chaos creates finer and finer structures, eventually resulting in a fractal. In an example of chaos in the real world, at a small enough scale, your model will stop being a good description of reality. Thermal noise and quantum effects will force you to change your model if nothing else has first. Models of chaos typically do have a largest scale. If you zoom far enough away from a strange attractor (for example), the motion looks like a dot twitching slightly. It is possible to average over this motion and get predictable results.

This can perhaps be best seen with the weather.[33] If you average the temperature over the next minute, how different will that be from the average temperature over the following minute? If you average the temperature over the next decade, how different will that be from the average temperature over the following decade? The averages are measured from temperature data or various paleoclimate proxies. Obvious periodic changes in temperature over the course of a day and a year are removed to only look at the chaotic fluctuations.[21] For times scales between 1 second and 10 days, as you average over longer times, the fluctuations get larger. It becomes harder to make predictions the more you average. Weather on these time scales is highly turbulent. For time scales between 10 days[22] and 10 years, the fluctuations get smaller as you average over longer times. This is getting less chaotic and taking averages allows you to make better predictions. For a weather forecaster, averaging over a decade makes much more sense than averaging over a week.

Many of the same techniques can be used both to predict a particular trajectory and to calculate averages of the motion. I will discuss several of these techniques, ordered roughly in increasing amount of knowledge about the dynamics required to use this technique: experimental analogues, direct numerical simulations, current machine learning techniques, using invariant quantities, mesoscale regularity, and periodic orbit theory.

7.1 Experimental Analogues

Sometimes, the best way to make a prediction is to measure what happens.[23] This allows you to see exactly what will happen in a particular instance. It does not allow you to predict exactly what will happen in other instances, unless you can return the system to the exact same initial conditions. This reinitialization is typically impossible in real chaotic systems, because it requires infinite precision for the initial conditions, but easy to do in computer models of chaos. Building an experimental analogue is also effective for measuring the statistical properties of the motion and exploring the space of what might happen.

The most well known of these analogous experiments are wind tunnels. For many situations involving fluids, the fluid will behave the same way if you make the size smaller and the velocity larger. Sometimes, you also have to change the density of the fluid, the temperature, or something else to get an analogous experiment.[24] This allows you to do tests using a scaled-down model with realistic fluid flow around it. Using wind tunnels to characterize fluid flow around an object is a standard practice in industry.[25]

Experimental analogues are less commonly used outside of wind tunnels because, for many chaotic systems, it is not possible to build an analogous scaled-down version. If the prediction is just as difficult to make as a prototype, there is no point in separating predictions and trials.

7.2 Direct Numerical Simulations

The equations of motion are typically differential or difference equations. If the model is a difference equation, you can pick an initial condition and evaluate the function numerically for each time step. Similarly, there are standard techniques for numerically solving differential equations, like Euler's method and the Runge-Kutta

[21] The figure described can be seen here: [33].

[22] Ten days is roughly the time is takes for weather systems to travel all the way around the earth, so averaging over longer time scales than this is like averaging over the entire atmosphere.

[23] This is sometimes called a postdiction.

[24] The goal is to make the Reynolds number, and all other relevant dimensionless parameters, the same in the original and scaled down versions.

[25] The Library of Congress had records of 89 wind tunnels in the United States in 2008.[34]

methods. These give a numerical approximation of the trajectory. This seems to be the most common way people try to predict chaos.

Direct numerical simulations work well for short-time predictions: up to about the Lyapunov time. For example, weather prediction uses direct numerical simulations and works pretty well for a few days. Trying to use direct numerical simulations to predict the trajectory for longer times does not work. Climate models built the same way as weather models are much less trustworthy.

Direct numerical simulations can still be useful for longer times if you look at statistics instead of individual trajectories. Although you cannot predict which trajectories will be followed, you sometimes can predict the statistics for a typical trajectory.

Direct numerical simulations are also useful to make pictures of what the chaos is doing and to get intuition about it. Many of the pictures in this report were created using direct numerical simulations.

7.3 Current Machine Learning

Machine learning has also been used to make predictions of chaotic systems.[35] Some methods are effective at both predicting the short-time trajectories,[26] up to a few Lyapunov times for sufficiently precise initial conditions, and for predicting the long-time statistics, including the tails of the probability distributions.[36]

A particularly interesting paper looking at the capabilities of artificial intelligence when dealing with chaos is *Long-term prediction of chaotic systems with machine learning* by Fan et al.[37] What they have done is interesting, but much less impressive than the title sounds. To achieve a 'long-term prediction,' they make a measurement of the chaotic system every few Lyapunov times and use the measure to update their 'prediction.' This does not allow you to know the long-time behavior in advance and is instead better thought of as a sequence of short-time predictions.

These methods can be trained on any data that looks chaotic: you do not need to have a good model for the equations of motion.[38] Machine learning can also be used to infer the values of unmeasured variables from a time series and to construct the equations of motion for a time series, although the variables and equations of motion produced are often not the most intuitive ones.[39] For example, if the observed data is position vs time, the most intuitive variable to add is the momentum. The machine learning algorithm might instead add some combination of position and momentum as its additional variable. These variables could also be used to fully characterize the motion, but they make it harder to see how this particular example connects to our broader understanding of classical mechanics.

7.4 Invariant Quantities

Energy is conserved in an isolated system. This is true regardless of whether there is chaos. There is at least one good prediction that we can make: the total amount of energy in the future will be the same as the total amount of energy now. There can also be other conserved quantities for some systems, like momentum, angular momentum, phase space area, or something more exotic.

The conserved quantities add constraints to the motion and confine it to a lower dimensional subspace of the state space: If energy is conserved, then the motion can only explore the 'surface' along which energy is constant, not the entire state space. The effects of these constraints can be dramatic: Hamiltonian chaos, which has conserved quantities, looks very different from a strange attractor, which does not.

Noether's Theorem tells us that there is an intrinsic connection between conserved quantities and symmetries.[40] If your chaotic system has a symmetry then, for every trajectory, you can find a corresponding trajectory symmetrically arrayed. It is often more important for recurrence to return to near a symmetrically equivalent location instead of returning to near the exact same location it started. Symmetries can be both a powerful tool for discovering the structure of the dynamics and a nuisance to do calculations with.

[26] Reservoir computing was the most effective of the three machine learning methods tried, which made accurate predictions of the results of a direct numerical simulation for 10 Lyapunov times.

For example, a hurricane is a dynamically important structure which can persist for weeks in the atmosphere. The motion of the winds around the hurricane is, on average, periodic. Except this motion is not really periodic because the center of the hurricane moves. As long as the hurricane is over warm open water, it does not matter where exactly the center of the hurricane is: the hurricane is symmetric under translations of the center of the hurricane. The winds around the hurricane are on average periodic, relative to the translation symmetry of the center of the hurricane.

Direct numerical simulations typically do not respect the conserved quantities and symmetries of the original motion. The small numerical errors might add or remove energy from the system. The calculated numerical trajectory can be qualitatively different from the actual trajectory, especially after a long time. This can cause problems with using direct numerical simulations to predict the long-time statistics. To avoid this problem, it is possible to design structure preserving algorithms: numerical schemes where the discrete approximation also conserves energy – or also shares the differential geometry of phase space. One sophisticated example of this is GeomPIC, which is used for plasma physics.[41] It is also possible to make machine learning algorithms which respect the invariant quantities and symmetries.[42]

The opposite of a conserved quantity is a Lyapunov function.[27] A Lyapunov function is a function of the state space which monotonically increases or decreases along every trajectory. Lyapunov functions are uncommon, but are extremely useful if you find one. They are a reliable way to show that an equilibrium is stable or to show that motion is not chaotic: if the Lyapunov function is always increasing along a trajectory, then the trajectory cannot return to close to where it has been previously. One example is the entropy of an ideal gas, which is a Lyapunov function for the equations of motion governing the probability distribution of molecules in position and velocity.[28]

7.5 Mesoscale Regularity

A general physical system might have chaos on multiple different scales of space or time. I described this above when discussing when taking averages allows you to make better predictions. Averaging over all of the chaos allows you to reliably characterize all of the statistical properties of the long-time behavior of the chaos.

In general, instead of just having chaos at small scales and non-chaotic motion (or no motion) at large scales, there could be more complicated combinations. There might be chaos at a small scale, non-chaotic motion at an intermediate mesoscale, chaotic motion at a larger scale, and perhaps more different behavior as you zoom in or out.

To what extent would the microscopic chaos affect the macroscopic chaos if there is a non-chaotic mesoscale in between?

When there is a non-chaotic mesoscale, then averages over the microscopic chaos will behave predictably. The macroscopic chaos will then depend on the mesoscopic averages, rather than on the microscopic chaos directly. There will be some fluctuations around these averages, but they will be small compared to the macroscopic chaos, especially if the mesoscale extends over several orders of magnitude.

One of the key lessons of chaos theory is that something should not be ignored simply because it is small. If there were no other sources of uncertainty, eventually this will be amplified and make the macroscopic motion inherently unpredictable. There are other sources of uncertainty, which are often larger and make the motion unpredictable sooner than the uncertainty provided by the macroscopic chaos: uncertainty in measuring the initial conditions, in unaccounted for external forces, or in what processes are relevant here. The microscopic chaos might not be the relevant source of uncertainty here.

If there is no non-chaotic mesoscale and there is chaos all the way down, then taking intermediate averages will not allow you to make better predictions. Thermal fluctuations and quantum uncertainties from within

[27] This is named after the same person as Lyapunov exponents, although it is a different thing.

[28] The equations of motion are called the Boltzmann equation, which describes how the probability of finding a particle at position x and with velocity v changes in time due to the particles' motion and collisions between particles. This probability converges to the Maxwell-Boltzmann distribution, as guaranteed by the H-Theorem, which proves it by showing that the entropy (H) is a Lyapunov function.

the system get amplified to the macroscopic scale.

Fluids often have this sort of mesoscale regularity. At a large scale, fluids are often turbulent. At a small scale, there is the random motion of atoms called molecular chaos. In between, viscosity dominates and the flow is not chaotic. In conditions easily achievable on the surface of the earth, this viscous mesoscale extends from the mean free path of a molecule (about $10^{-10}m$ for water or $10^{-7}m$ for air) to about a millimeter ($10^{-3}m$). It is reasonable to describe fluid turbulence using only mesoscopic averages like pressure and temperature. Molecular chaos does cause fluctuations in these average quantities, but the fluctuations are many orders of magnitude smaller than the smallest turbulent vortex. The most important source of uncertainty for turbulence in macroscopic fluids is not the microscopic motion of molecules. Instead, small magnitude but large scale fluctuations are more important, like changes in the amount of light produced by the sun or wobbles of the earth's axis of rotation. Some of these external uncertainties might be chaotic and impossible to predict themselves, while others are not.

A computer simulation of a chaotic system can be thought of in a similar manner. At the smallest scales, atoms move randomly. However, computer engineers are careful to not allow this motion to affect the state of the computer. Bits are designed to only flip when they are supposed to, not as a result of thermal noise or anything else. When you run a simulation of a chaotic system on a computer, the macroscopic behavior is chaotic, but it is shielded from the atomic motion by the efforts of the computer engineers. Models of chaos run on a computer are easy to reinitialize or copy exactly, because it is cut off from microscopic motion that is hard to measure or control.

7.6 Periodic Orbit Theory

The most reliable way to average over chaotic motion, which provably converges to any statistical property (average, variance, correlations, etc.) of the long-time motion, is periodic orbit theory.

While it is possible to estimate the statistical properties of a chaotic system using direct numerical simulation, machine learning, or experimental analogues, these techniques tend to have difficulty estimating small scale structure and low probability events without large numbers of experiments or runs. The probability distributions describing where an object might be in a chaotic system are often very spiky and have weird tails. It is not uncommon for the distribution to only be nonzero on a fractal set,[29] and for the distribution along this set to be fractal as well.[30] This makes the statistics extremely hard to characterize using conventional numerical methods. To get a feel for the sorts of things that these distributions can do, a numerical approximation for the Hénon map is shown in Figure 9.

There is another way, called periodic orbit theory, to get better statistical knowledge of the dynamics, if the chaos is ergodic or has a strange attractor which contains dense periodic orbits (or relative periodic orbits if there is symmetry).

It is also helpful to assume that the chaos is entirely connected. If it is not, because there are island chains or because there are multiple strange attractors, then the chaos is multistable and the statistical predictions depend on which island chains you start between or which basin of attraction you start in.[31] Periodic orbit theory works less well for Hamiltonian chaos, where there are often many island chains, than it does when there is a single strange attractor.

Chaos often contains dense periodic orbits, as a result of recurrence and shadowing lemmas. A shadowing lemma proves that, under certain conditions, if you find a trajectory which approximately does something, then there exists a trajectory nearby which exactly does it. Because chaos is recurrent, the trajectory will return to close to itself after some time. The shadowing lemma guarantees that there exists a nearby initial condition which exactly returns to itself after a similar amount of time. This is a periodic orbit, close to any initial condition which is recurrent.

[29] If you plot the places where the distribution is not zero, the result is a fractal. A fractal is a shape which continues to have interesting structure as you zoom in on it.

[30] As you move along this set, the distribution itself rises and falls in a fractal manner, often looking like a Weierstrauss function.

[31] Recall the definitions of 'multistable' and 'basin of attractor' from Section 4.2.

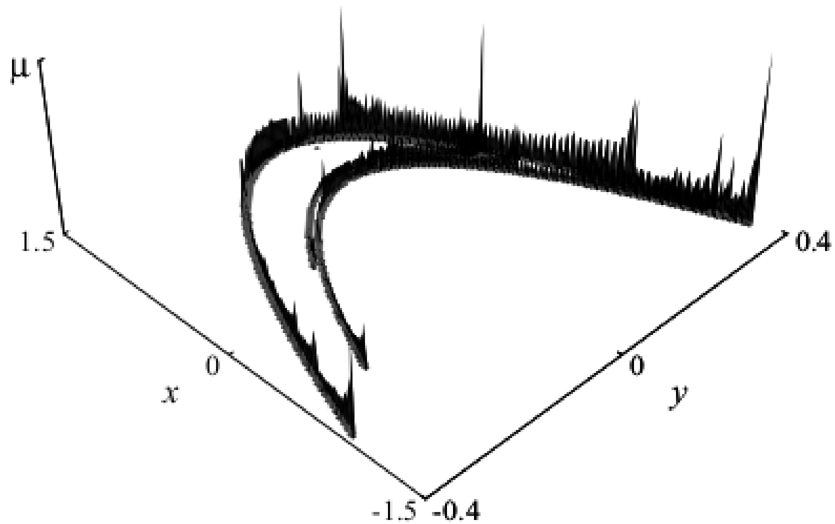

Figure 9: A numerical approximation of the probability distribution for motion along the Hénon attractor, showing the fractal nature of the distribution. Note that the numerical method likely does not converge, so the actual distribution is even messier than what is shown.[43]

We do not usually see these periodic orbits because they are all unstable. If you start at the exact right initial condition, then the motion will be periodic. If you start even slightly off, then the motion starts off following the periodic orbit, but then diverges away from it due to the sensitive dependence on initial conditions. The periodic orbits are a dense set (close to everywhere) of measure zero (almost all initial conditions are not periodic), like the rational numbers among the real numbers. For the shift map,[32] the periodic orbits are exactly the rationals among the reals.

Periodic orbits are great because they connect short-time and long-time motion. It takes a finite amount of precision to find a periodic orbit,[33] but once you have found one, you know what its motion will be from then on. Since the periodic orbits are dense, any motion in the chaos can be approximated for a short time by something periodic. You can think of a typical chaotic trajectory as starting near one periodic orbit, following it for a bit, then diverging away. It then follows a different periodic orbit a bit before diverging away again. The orbit goes through a sequence of neighborhoods of periodic orbits.

Making a statistical prediction about the long-time behavior of a typical chaotic trajectory can be converted into a statistical prediction involving the neighborhoods of periodic orbits. This allows you to gain long-time statistical knowledge from short-time calculations.

The calculation involves a sum over all periodic orbits, since the trajectory could move near any of them. In practice, you can only sum over a finite number of them. If the period of the shortest periodic orbit that is missing from your sum is τ, then the error scales as $e^{-\tau}$. This is fast convergence, and the convergence is not dramatically different for different parts of the distribution function, like the tails.

There are multiple statistical things you might want to predict (mean, standard deviation, etc.). Doing a sum over all periodic orbits is hard, so we only want to do it once for each chaotic system. Instead of directly calculating the average we want, we instead calculate something called a trace formula. The trace formula is not physically relevant itself, but it can be used to calculate any statistical prediction you want – much like the partition function of statistical mechanics. The end goal of periodic orbit theory is not the

[32]The shift map was described in Section 3.1.

[33]In order to reliably identify a trajectory of time T, you have to be confident that whatever errors are introduced by your identification method (often a direct numerical simulation) have not had enough time to grow large. This can be done, as long as you make sure that the initial uncertainty is small enough. For a typical chaotic trajectory, reliably describing its long time behavior requires exponentially increasing initial precision. For periodic motion, you only need enough precision for it to complete one period. After that, you know that it will continue to repeat what it has done before.

fractal distribution describing the long-time motion itself, but rather a smooth function which allows you to calculate any statistical information about the distribution.

The details of how to actually calculate a trace formula for a particular chaotic system can be found in ChaosBook.[43]

8 Control of Chaos

Control theory is a separate field of mathematics from chaos theory, which looks at to what extent it is possible to control the behavior of some dynamical system given a particular set of inputs to that system. For example, you might have a complicated circuit and you can only control the voltage at a few points of the circuit. How well can you control the voltages and currents elsewhere in the circuit?

A chaotic system can move in lots of different ways. You might want to control which way the chaos actually does move. You have some ways to regularly measure the system, and some ways to influence the system, in response to what you measure. Ideally, you use small inputs, not to overpower the chaos, but to continually guide the chaos into doing the things you want it to do.[34]

Several techniques have been developed to apply control theory to chaotic systems. The OGY (Ott, Grebogi, & Yorke) Method adjusts the orbit each time it passes through a Poincaré section.[44] The Pyragas Method uses a continuous input to control the entire path of the orbit.[45] The goal of both methods is to make the motion non-chaotic by forcing it to follow one of the unstable periodic orbits embedded in the chaos.

There are two things needed for chaotic control theory to work. (1) You have to know the periodic orbit ahead of time. In order to regularly or continually do small corrections to keep the orbit on track, you have to know what track you are keeping it on. (2) The inputs you control cannot be orthogonal to the directions in which the motion is unstable. For each periodic orbit, there are typically some ways of changing the initial conditions which cause the trajectories to separate exponentially (unstable directions) and some ways of changing the initial conditions which do not, or contract exponentially (stable directions). The unstable directions govern the 'stretch' and the stable directions govern the 'fold' or 'stretch and fold.' Controlling inputs in the stable directions is not useful. You need to have control over inputs in each of the unstable directions. As the actual trajectory starts to diverge from the orbit you want it to follow, add a small input to nudge the trajectory back on track.

Attempting to control fluid turbulence using a sub-millimeter fan will not work, even if it is a strong fan, because this is in the viscous scale and any perturbation will quickly damp away. You instead need something at a larger length scale, even if the thing you control at the larger length scale only has a small direct effect.

One example of the use of chaotic control theory is in spacecraft orbit design, especially when fuel concerns are more important than travel time. First, numerically approximate the three body problem, usually Earth - moon - spacecraft. Find an approximate trajectory which goes from close to the spacecraft's current location to close to where you want it to go. This is not a real trajectory, but it should be 'close' to a real trajectory in a local sense.[35] At every point along the trajectory, if you start close to this point, the subsequent motion will be close to the trajectory for a short time (but not for longer than the Lyapunov time). Continually track the location of the spacecraft. Whenever you observe that the spacecraft is deviating from the trajectory you want it to follow, use small thruster rockets to push it back on track. You will need to do lots of little thrusts: at least one per Lyapunov time. But the total amount of thrust needed is small, so this process requires very little fuel. One example of this is shown in Figure 10.

9 Other Predictions

When there is chaos, it is impossible to make predictions of a particular trajectory for much longer than the Lyapunov time. It is still possible to make statistical predictions, as long as you average over a scale

[34] A good introduction to Controlling Chaos can be found on Scholarpedia, written by one of the people who developed these methods.

[35] We could also use a shadowing lemma to show that there is a real trajectory nearby which does something similar.

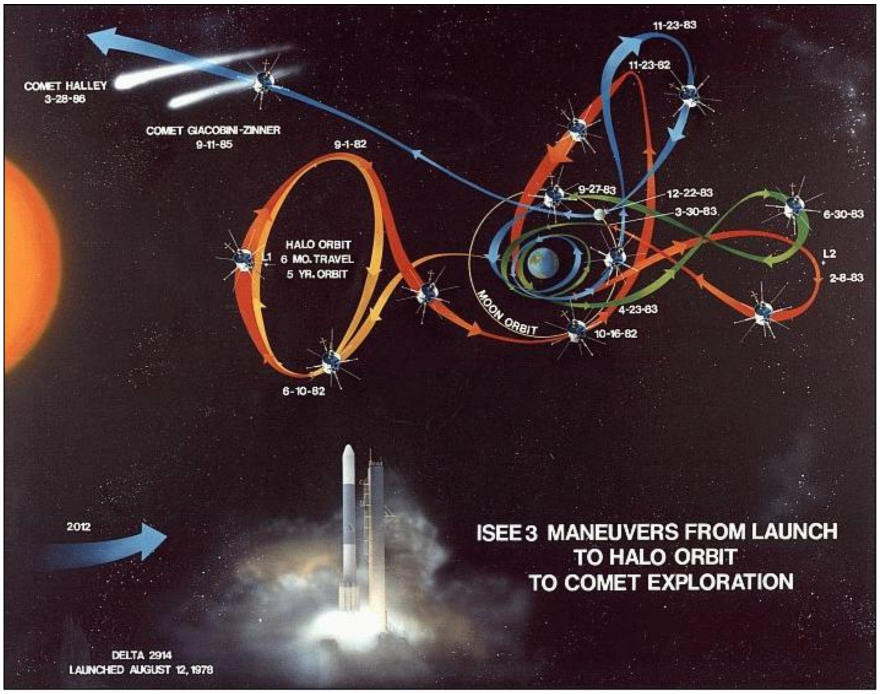

Figure 10: The International Sun-Earth Explorer-3 spacecraft was launched in 1978 to study the outer magnetosphere and solar wind. By 1982, it had fulfilled its primary mission. NASA & ESA decided that it would be the first spacecraft to visit a comet, renaming it the International Cometary Explorer. They found this approximate chaotic orbit that went from its current location, around the Earth and moon repeatedly, to pass through the tail of Comet Giacobini–Zinner in September 1985. Chaotic control theory was used to make the spacecraft follow this path, using the small amount of fuel remaining after it completed its original mission.[46]

larger than the chaos. If there is chaos on many scales, and you try to average at an intermediate scale, then the averages will themselves be unpredictable. Even this does not mean that there is nothing that can be predicted.

Finding additional things that can be predicted is fun. It shows us that there is universal structure, some of which is probably still to be discovered, even in the midst of chaos. But it is often hard to convert these predictions into something other people care about. What is predictable tends to be far removed from what you originally wanted to predict.

I will now briefly describe two examples of these types of predictions: turbulent spectra and the Feigenbaum constants.

9.1 Turbulent Spectra

For many fluids, their behavior is mostly independent of scale for a large range of scales. A cubic meter of water behaves in a similar way to a cubic kilometer of water.[36] This causes the turbulence to be self-similar. By thinking about how energy is transferred between length scales, we can predict a relationship between statistical properties at different scales.

The classic version of this argument is due to Kolmogorov.[47] He showed that for uniform, (statistically) stationary, three dimensional, isotropic turbulence, the energy contained in an eddy with size length L is

[36] Although it does not behave in a similar way to a cubic millimeter of water, because viscosity is much more important there.

proportional to $L^{5/3}$ (in terms of the wave number, k, this is $k^{-5/3}$). For other kinds of energy transfer, a turbulent spectrum can also be predicted, with a different exponent.

Experimental evidence for turbulent spectra is hard to get. In order to see it, there has to be multiple orders of magnitude between the scale at which energy is added and the scale at which viscosity dissipates the energy into heat. Measurements have to be precise enough at all of these scales to get a clear fit to the power law. Kolmogorov-like spectra have been observed in the ocean,[48] but the best evidence for turbulent spectra come from spacecraft measurements of the solar wind. The energy is input at a very large scale (the sun) and viscosity is small because the density is so low. There are multiple energy transfer mechanisms that are important at different scales so you can see multiple power laws with different exponents.[49]

If the structure of your system is very different on different lengths scales, as in most biological systems, then there is no reason for the motion to be self-similar or for there to be predictable relationships between statistical properties of the motion at different scales. Predictable turbulent spectra typically do not occur in these systems.

9.2 Feigenbaum Constants

Feigenbaum constants characterize one of the most common ways chaos can arise from non-chaotic motion. The constants were originally derived using the logistic map, but they have been shown to be universal for any period doubling cascade.

A periodic doubling cascade is a process by which something which starts out not moving begins to move periodically, and then chaotically. Start with something with a stable equilibrium: regardless of where you start, the motion will settle down into something stationary. Continually add energy to it (or change the parameters of the system in some other way) so that it starts to move periodically. Increasing the amount of energy input will typically cause the amplitude of the periodic motion to increase. For many systems, at some point, the period of the periodic motion will double. Instead of returning to where it was before each time around, it returns to a slightly different place after one time around and only repeats after two times around. Increasing the energy input even more will often cause the period of the oscillation to double again: the motion returns to where it was because only after four times around. If these period doublings continue to happen as you increase the energy input, then you have found a period doubling cascade.

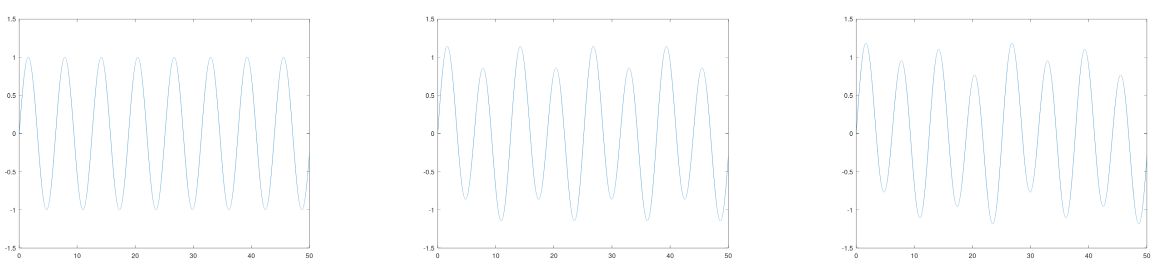

Figure 11: (left) Simple periodic motion. (middle) Periodic motion with twice the period. (right) Periodic motion with four times the period. These were made by adding sine waves of different wavelengths together, not by a dynamical process undergoing a period doubling cascade. *Made by Jeffrey Heninger.*

The amount of additional energy input needed to cause the next period doubling decreases. It takes more additional energy input to go from being stationary to moving periodically than it takes to go from moving periodically to moving moving periodically in a more complicated way. These transitions get closer and closer together until there is a limit point, beyond which the motion is chaotic. It is as though the period of the motion doubled again and again until it reached infinity, at which point the motion becomes chaotic instead of periodic.

Call the amount of energy input when the stationary system first becomes periodic r_0, the amount of energy input when the periodic motion doubles its period the first time r_1, the amount of energy input when the periodic motion doubles its period the second time r_2, etc. I said that the differences between these continue to get smaller, so $\lim_{n\to\infty} r_n - r_{n-1} = 0$. We can also look at the amount that the interval difference shrinks by between each pair of period doublings. This approaches a constant:

$$\lim_{n\to\infty} \frac{r_{n-1} - r_{n-2}}{r_n - r_{n-1}} = \delta$$

This is the first Feigenbaum constant, and it is the same for every period doubling cascade: $\delta = 4.669\ldots$. It has been measured in the logistic map,[50] convection in water[51] and mercury,[52] electronic nonlinear oscillators,[53] Josephson junctions between superconductors,[54] and other physical systems.

The period doubling cascade does not depend on the details of what exactly you're looking at as you get close to the transition to chaos. Instead, the transition to chaos is universal. And it is not just one constant: at the limit point of the period doubling cascade, the motion approaches a particular Cantor set, characterized by the Feigenbaum constants.

Chaos in many physical systems arises through the same process, which leads to the exact same mathematical structure at the point where periodic motion transitions into chaos. Whenever you see something that looks like the start of a period doubling cascade, you could make a surprisingly precise prediction: this particular limit will approach this particular number. It is not clear how this prediction would allow you to cause a significant impact on the world. There exists interesting structure within chaos which allow you to make some predictions, but it does not allow you to make most of the predictions you want about the future.

10 Conclusion

It is surprising how long it took to develop chaos theory. While some of the foundations were laid down by Poincaré and Lyapunov in the late 1800s, most of what we know about chaos has been discovered since 1950. Some physical systems we now know are chaotic had been thought about or experimented on for hundreds of years previously, like the three body problem of celestial mechanics. What seems to have happened is that people did not recognize chaos without a category to describe it. When an experiment would behave erratically, people would have considered it a failed experiment, instead of a successful experiment of erratic behavior.

Engineers still try to avoid chaos. Designing something is much easier if what you designed behaves in a predictable way. This creates a bias towards designed things which are not chaotic. When interactions with fluids make the chaos unavoidable, engineers tend to build a model and put it in a wind tunnel to see what empirically happens.

Once you know what chaos is, it is not difficult to recognize it. Recurrence is easy to check, so the main question is whether the motion has a positive Lyapunov exponent. Do trajectories which begin close together diverge exponentially? There are straightforward ways to determine this either for equations of motion or an experimental time series.

When nearby trajectories diverge, any initial uncertainty grows exponentially until it becomes impossible to predict what particular path will be followed. It is often still possible to determine the statistical properties of the motion, using experimental analogues, direct numerical simulations, current machine learning techniques, or periodic orbit theory. Knowledge of the invariant quantities and symmetries is often helpful.

When calculating the statistical properties, you should calculate the averages on a scale larger than the chaos. Otherwise, the averages themselves can change chaotically. It is often the case that there will be some large or intermediate scale which is not chaotic, and you can average over the chaotic or random motion on smaller scales in order to describe the macroscopic motion. But some times, all scales between the atomic scale and the scale you are interested in are chaotic, which allows the intrinsic randomness of the atomic scale to be amplified all the way up to the macroscopic.

Even when chaos renders both particular trajectories and averages unpredictable, that does not mean that there is nothing that can be predicted. Chaos theory has shown that there sometimes are surprising patterns within the chaos, even if they are not the predictions you originally hoped to make.

It is also possible sometimes to control chaos. In order for you to be able to control some chaotic motion, you have to have inputs to the motion for each variable which is unstable, and you have to be able to measure and react to the motion faster than the Lyapunov time. For sufficiently complicated or difficult to access physical systems, this might not be possible.

Chaos renders most of the details about the future unpredictable. The exponential growth in uncertainty for a chaotic system quickly wipes out most of the information you started with. But not everything about the motion is unpredictable. Sometimes the statistics of the motion are predictable. Sometimes there are conserved quantities. Sometimes there are other more surprising patterns. Sometimes it is possible to control the chaos. But these situations are not guaranteed. The statistics can be non-stationary. The conserved quantities can be lost to the environment. The control system might requite extremely high resolution in space or time in places which you cannot access. In these situations, the future is truly unpredictable.

References

[1] Heninger & Johnson. *Chaos in Humans.* AI Impacts. (2023) https://wiki.aiimpacts.org/doku.php?id=uncategorized:ai_safe

[2] Heninger & Johnson. *AI Safety Arguments Affected by Chaos.* AI Impacts. (2023) https://wiki.aiimpacts.org/doku.php?id=uncategorized:ai_safety_arguments_affected_by_chaos.

[3] Heninger & Johnson. *Superintelligence Is Not Omniscience.* AI Impacts. (2023) https://blog.aiimpacts.org/p/superintelligence-is-not-omniscience.

[4] Heninger. *You Can't Predict A Game of Pinball.* AI Impacts. (2023) https://blog.aiimpacts.org/p/you-cant-predict-a-game-of-pinball.

[5] Heninger. *Whole Bird Emulation Requires Quantum Mechanics.* AI Impacts. (2023) https://blog.aiimpacts.org/p/whole-bird-emulation-requires-quantum-mechanics.

[6] Lorenz. Quoted by Danforth. *Chaos in an Atmosphere Hanging on a Wall.* Mathematics of Planet Earth. (2006) http://mpe.dimacs.rutgers.edu/2013/03/17/chaos-in-an-atmosphere-hanging-on-a-wall/.

[7] Wolfram. *A New Kind of Science.* (2002) p. 739. https://www.wolframscience.com/nks/p739–computational-irreducibility/.

[8] Wolfram MathWorld. *Computational Irreducibility.* (Accessed Jan 11, 2023) https://mathworld.wolfram.com/ComputationalIrreducibility.html.

[9] Wolfram MathWorld. *Principle of Computational Equivalence.* (Accessed Jan 11, 2023) https://mathworld.wolfram.com/PrincipleofComputationalEquivalence.html.

[10] Israeli & Goldenfeld. *Coarse-graining of cellular automata, emergence, and the predictability of complex systems.* Physical Review E **73**. (2005) https://arxiv.org/pdf/nlin/0508033.pdf.

[11] Eliezer Yudkowsky. *Markets are Anti-Inductive.* LessWrong. (2009) https://www.lesswrong.com/posts/h24JGbmweNpWZfBkM/markets-are-anti-inductive.

[12] Cross. *The Behavior of Stock Prices on Fridays and Mondays.* Financial Analysts Journal **29.6**. (1973) p. 67-69. https://www.jstor.org/stable/4529641#metadata_info_tab_contents.

[13] Scott Alexander. *The Phatic and the Anti-Inductive.* Slate Star Codex. (2015) https://slatestarcodex.com/2015/01/11/the-phatic-and-the-anti-inductive/.

[14] Huxley. *The Anniversary Address of the President.* The Quarterly Journal of the Geological Society of London **25**. (1869) p. xlv. https://www.google.com/books/edition/The_Quarterly_Journal_of_the_Geological.

[15] Maxwell. *On the Dynamical Theory of Gases.* Philosophical Transactions of the Royal Society of London **157**. (1867) p. 49-88. https://royalsocietypublishing.org/doi/pdf/10.1098/rstl.1867.0004.

[16] Strogatz. *Nonlinear Dynamics And Chaos: With Applications To Physics, Biology, Chemistry, and Engineering.* (CRC Press, 2000).

[17] Chirikov. *Resonance Processes in Magnetic Traps.* Journal of Nuclear Energy, Part C: Plasma Physics **1**. (1960) p. 253-260. https://www.quantware.ups-tlse.fr/chirikov/refs/chi1959e.pdf

[18] Keil. *Stability of Particle Orbits.* Classical Mechanics. (1988) https://cds.cern.ch/record/187614/files/cer-000099059.pdf

[19] Petrosky. *Chaos and Cometary Clouds in the Solar System.* Physics Letters A **117.7**. (1986) p. 328-332. https://www.sciencedirect.com/science/article/abs/pii/0375960186906730

[20] Casati et al. *Hydrogen Atom in Monochromatic Field: Chaos and Dynamical Photonic Localization.* IEEE Journal of Quantum Electronics **24.7**. (1988) p. 1420-1444. https://www.researchgate.net/profile/Giulio-Casati/publication/299683556_Relevance_of_Chaos_in_Quantum_Mechanics_The_Hydrogen_Atom_in_the_Microwave_Field/ of-Chaos-in-Quantum-Mechanics-The-Hydrogen-Atom-in-the-Microwave-Field.pdf

[21] Leonardo da Vinci. *Studies of Water.* (c. 1510-1512) RCIN 912660.

[22] Sharkovskii. *Coexistence of cycles of a continuous map of a line into itself.* Ukrainian Mathematical Journal **16**. (1964) p. 61-71.

[23] Li & Yorke. *Period Three Implies Chaos.* The American Mathematical Monthly **82.10**. (1975) p. 985-992. https://yorke.umd.edu/papers/Li-Yorke%20Period%20Three%20Implies%20Chaos.pdf.

[24] Lorenz. *Deterministic nonperiodic flow.* Journal of the Atmospheric Sciences **20**. (1963) p. 130-141. https://journals.ametsoc.org/view/journals/atsc/20/2/1520-0469_1963_020_0130_dnf_2_0_co_2.xml?tab_body=pdf.

[25] Kolmogorov. *On the Conservation of Conditionally Periodic Motions under Small Perturbation of the Hamiltonian.* Proceedings of the USSR Academy of Sciences **98**. (1954)

[26] Moser. *On invariant curves of area-preserving mappings of an annulus.* Nachrichten der Akademie der Wissenschaften zu Göttingen II. (1962)

[27] Arnold. *Proof of a theorem of A. N. Kolmogorov on the preservation of conditionally periodic motions under a small perturbation of the Hamiltonian.* Russian Mathematical Surveys **18.5**. (1963) p. 9-36.

[28] Wolf et al. *Determining Lyapunov Exponents from a Time Series.* Physica D: Nonlinear Phenomena **16.3**. (1985) https://hal.archives-ouvertes.fr/hal-01654059/document.

[29] Noll, Mende, & Sisodiya. *Method for seeding a pseudo-random number generator with a cryptographic hash of a digitization of a chaotic system.* Patent US5732138. (1996) https://patents.google.com/patent/US5732138.

[30] Lan et al. *Testing and Prediction of Traffic Flow Dynamics with Chaos.* Journal of the Eastern Asia Society for Transportation Studies **5**. (2003) https://www.researchgate.net/profile/Lawrence-Lan-2/publication/266208103_Testing_and_Prediction_of_Traffic_Flow_Dynamics_with_Chaos/links/542a76e10cf27e39fa8e9648/ and-Prediction-of-Traffic-Flow-Dynamics-with-Chaos.pdf.

[31] Nagatani. *Chaotic jam and phase transition in traffic flow with passing.* Physical Review E **60**. (1999) https://journals.aps.org/pre/abstract/10.1103/PhysRevE.60.1535.

[32] Disbro & Frame. *Traffic flow theory and chaotic behavior.* Engineering Research and Development Bureau: New York State Department of Transportation: Special Report 91. (1989) https://rosap.ntl.bts.gov/view/dot/15604.

[33] Lovejoy. *Weather, Macroweather, and the Climate.* (Oxford University Press, 2019) Figure 2.4.
See also: Lovejoy. *Spectra, Intermittency, and Extremes of Weather, Macroweather, and Climate.* Scientific Reports **8**. (2018) Figure 1. https://www.nature.com/articles/s41598-018-30829-4.

[34] Goodrich & Gorham. *Wind Tunnels of the Western Hemisphere.* Federal Research Division, Library of Congress. (2008) p. vii. https://loc.gov/rr//frd/pdf-files/Western_Hemisphere_Wind_Tunnels.pdf

[35] Tang et al. *Introduction to Focus Issue: When machine learning meets complex systems: Networks, chaos, and nonlinear dynamics.* Chaos **30**. (2020) https://www.pik-potsdam.de/members/kurths/publikationen/2020/tang-2020-5-0016505.pdf.

[36] Chattopadhyay et al. *Data-driven predictions of a multiscale Lorenz96 chaotic system using machine-learning methods: reservoir computing, artificial neural net, and long short-term memory net.* Nonlinear Processes in Geophysics **27.3**. (2020) https://npg.copernicus.org/articles/27/373/2020/.

[37] Fan et al. *Long-term prediction of chaotic systems with machine learning.* Physical Review Research **2**. (2020) https://journals.aps.org/prresearch/pdf/10.1103/PhysRevResearch.2.012080.

[38] Pathak et al. *Model-Free Prediction of Large Spatiotemporally Chaotic Systems from Data: A Reservoir Computing Approach.* Physical Review Letters **120**. (2018) https://www.lri.fr/~gcharpia/deeppractice/2022/chap_5_biblio/Chaos/PhysRevLett.pdf.

[39] Lu et al. *Attractor Reconstruction by Machine Learning.* Chaos: An Interdisciplinary Journal of Nonlinear Science **28**. (2018) https://arxiv.org/pdf/1805.03362.pdf.

[40] Noether. *Invariante Variationsprobleme.* Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse. (1918) p 235-257.
See also Wikipedia: Noether's Theorem.

[41] *Geometric Numerical Integration and Reduced Complexity Modelling.* Max Planck Institute for Plasma Physics. https://www.ipp.mpg.de/4105128/gspm.

[42] Mattheakis et al. *Hamiltonian neural networks for solving equations of motion.* Physical Review E **105**. (2022) https://arxiv.org/pdf/2001.11107.pdf.

[43] Cvitanović et al. *Chaos: Classical and Quantum.* Ch. 19: Transforming Densities. https://chaosbook.org/.

[44] Ott et al. *Controlling chaos.* Physical Review Letters **64.11**. (1990) https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.64.1196.

[45] Pyragas. *Continuous control of chaos by self-controlling feedback.* Physics Letters A **170.6**. (1992) p. 421-428. https://www.sciencedirect.com/science/article/abs/pii/0375960192907458.

[46] eoPortal: Satellite Missions Catalogue. *ISEE-3 / ICE.* (2014) Accessed 2023-01-30. https://www.eoportal.org/satellite-missions/isee-3#mission-status.

[47] Kolmogorov. *Dissipation of energy in the locally isotropic turbulence.* Proceedings of the USSR Academy of Sciences. (1941)

[48] Grant et al. *Turbulence spectra from a tidal channel.* Journal of Fluid Mechanics 12.2. (1962) https://www.damtp.cam.ac.uk/user/tong/fluids/vancouver.pdf.

[49] Kiyani et al. *Dissipation and heating in solar wind turbulence: from the macro to the micro and back again.* Philosophical Transactions of the Royal Society A **373**. (2015) https://royalsocietypublishing.org/doi/full/10.1098/rsta.2014.0155.

[50] Feigenbaum. *Quantitative universality for a class of nonlinear transformations.* Journal of Statistical Physics **19**. (1978) p. 25-52. https://link.springer.com/article/10.1007/BF01020332.

[51] Giglio et al. *Transition to Chaotic Behavior via a Reproducible Sequence of Period-Doubling Bifurcations.* Physical Review Letters **47.4**. (1981) https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.47.243.

[52] Libchaber et al. *Period doubling cascade in mercury, a quantitative measurement.* Journal de Physique Lettres **43.7**. (1982) https://hal.science/file/index/docid/232033/filename/ajp-jphyslet_1982_43_7_211_0.pdf.

[53] Linsay. *Period Doubling and Chaotic Behavior in a Driven Anharmonic Oscillator.* Physical Review Letters **47.19**. (1981) https://www.phys.hawaii.edu/phys480l/wp-content/uploads/sites/9/2019/08/Linsay_chaos.pdf.

[54] Yeh & Kao. *Universal Scaling and Chaotic Behavior of a Josephson-Junction Analog.* Physical Review Letters **49.27**. (1982) https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.49.1888.