

APPLIED DYNAMICAL SYSTEMS

1st edition DRAFT

Copyright © 2025 Robert Ghrist
All rights reserved worldwide
Published by Agenbyte Press
Jenkintown PA, USA

ISBN 978-1-944655-XX-X



*Robert Ghrist is the Andrea Mitchell University PIK Professor
of Mathematics and Electrical & Systems Engineering
at the University of Pennsylvania*

PREFACE



THIS TEXT covers dynamical systems in continuous & discrete time. It serves as a self-contained introduction to linear algebra (assuming familiarity with matrix algebra), using systems modelling as motivation.

The book is intended for undergraduate students in Engineering & the Natural Sciences. As such, less emphasis is placed on careful proofs of major theorems; more emphasis is placed on intuition, applications, and the development of systems-level thinking. Students looking for much more or much less rigor will find no end of texts clustered at the extremes. This book is a balance.

The integration of linear algebra in the context of systems modelling and dynamics is an unusual choice, meant to soften the abstraction that sometimes leaves students wondering at the rationale for eigenvalues and diagonalization.

Several unique features should appeal to a broad audience of learners: (1) The marriage of continuous and discrete dynamics is effected with the use of continuous and discrete calculus; (2) The nature and centrality of eigenvalues and eigenvectors is integral to the story; (3) The mathematical aspects of the theory are front-and-center, especially the many and various ways that ideas from geometry and topology arise; (4) There are many figures, pedagogical even when fanciful.


The text is ideally paired with video, and the animations of concepts & examples are a fundamental component of learning the material: the mathematics of evolution is not easily borne on static images.

The author wishes to thank his many students in the School of Engineering and Applied Sciences at the University of Pennsylvania. In addition, during the writing of this text, the author's research has received generous support from the Assistant Secretary of Defense for Research & Engineering, the Air Force Research Labs, DARPA, the National Science Foundation, and Andrea Mitchell.

For Philip Holmes
Applied Mathematician : Pure Poet
il miglior fabbro.

VOLUME I

THE FIRST DIMENSION




CHAPTER 1 : DYNAMICS & BIFURCATIONS

DYNAMICAL SYSTEMS is the study of behaviors of systems that change over time. Growth, decay, oscillation, evolution, collapse, and chaos are all examples of behaviors that systems exhibit. If Mathematics is the Science of pattern, then dynamical systems is the mathematics of behavior.

TIME: CONTINUOUS & DISCRETE

Systems evolve over time. The fundamental dichotomy of this text lies in the representation of time as continuous or discrete. Continuous-time dynamical systems are described at the infinitesimal level using vector fields or differential equations. The smooth clockwork-like motion of heavenly bodies is the great example of continuous-time dynamics. Discrete-time systems can arise through the discrete nature of data – perhaps from economic data, which are collected only on a periodic basis. Other examples of discrete-time dynamics come from sampling a system at certain events, as opposed to sampling at uniform intervals of time.


EXAMPLE: A ball bouncing on a table leads to a discrete-time dynamical system, where one can record the magnitude of velocity of successive impacts. This would have very simple dynamics, as the impact velocities would (under natural physical assumptions) decrease exponentially over time, converging to zero. Equivalently, one could record instead the elapsed time between bounces: this too would lead to a monotone decrease converging to zero. The behavior of this system increases greatly in richness when the table is being oscillated




vertically, feeding energy into the system.

STATE SPACE

The dynamics of a system – its evolution – is specified by tracking the state of a system and how it changes over time. For some systems, the state is easily quantified as a function of time – think the price of crude oil or the number of cells in a tumor. Other systems are more subtle and require care in the specification of state. A rigid-rod pendulum has state given by angle of the rod (say, as measured from the vertical) and angular velocity. Additional features (such as angular acceleration) could be included in the state but are redundant (acceleration determines velocity up to an initial condition). A great many systems have multiple equivalent descriptions or coordinatizations of state space.



As indicated by the name, a **state space** (frequently and confusingly called a **phase space** by physicists) is a topological space of states, usually outfitted with



geometry and explicit coordinates. Such exact descriptors do not obviate the need for awareness of topological issues, however. For example, the angular velocity v of a rigid-rod pendulum takes values in the reals, $\mathbb{R}$, where the sign determines increase or decrease of angle (subject to a chosen orientation); the angular variable θ takes values in a circle, denoted $\mathbb{S}^1$, commonly coordinatized in radians as the reals modulo 2π . The space of all possible ordered pairs (θ, v) is the state space formally denoted $\mathbb{S}^1 \times \mathbb{R}$ and visualized as either


an annulus or a cylinder (without ends). For most applications in this text, state spaces are $\mathbb{R}^n$, the n -dimensional Euclidean space of ordered n -tuples of reals. Occasionally $\mathbb{S}^1$ will make an appearance. The possibilities for state spaces in general are manifold; we will not dwell on such.

The key feature of an appropriate state space for a system is that the dynamics are purely a function of state: knowing where you are now tells you where to go next. If one wants to consider dynamics that are nonautonomous – dynamics that change over time – then it is appropriate to include time within the state space.

QUANTITATIVE VS. QUALITATIVE

Our first major intellectual leap will be a departure from quantitative specification of complete system behaviors to qualitative classification of features. This is a flight of necessity, as nonlinear systems, even if deterministic, do not typically admit explicit solutions expressible in terms of the “nice” functions from calculus. For example, a rigid-rod pendulum swinging back-and-forth has a simple differential equation model whose exact solutions require the use of elliptic integrals – well outside the bounds of the typical calculus student experience. Slightly more complex models either use greatly more complex technology (e.g., hypergeometric functions) or else have no exact solutions expressible in known functions. Fortunately, exact solutions are not always or often necessary.

EXAMPLE: Consider a rigid-rod pendulum suspended over a floor with two strong magnets, one on either side of the bob’s minimum, each attracting and pulling the (magnetic) bob towards it. To model this system would require knowing the magnetic force, the precise nature of the attraction, the exact locations of the magnets, and more. The resulting differential equation would likely be unpleasant, nonlinear, and unsolvable. Yet, if you give the pendulum a gentle swing, what happens in clear. Eventually, it comes to rest at one of two states, pulled to one of the magnets. From that resting state, a small tap quickly dissipates. In such a system one can classify qualitative long-term behavior without knowing exact quantitative details.




EQUILIBRIA

The core dynamical feature on which behaviors are built is an equilibrium state. An **equilibrium** is a point in the state space which does not change under the dynamics over time: stationary solutions or fixed points are synonyms for equilibria.

EXAMPLE: A rigid-rod pendulum, whether damped or frictionless, has an obvious equilibrium solution in which the bob hangs vertically below with no angular velocity.

EXAMPLE: In Machine Learning, a typical neural network has a large state space of weights. The goal is to tune the weights as to minimize a cost functional, say Φ , that indicates the error of the system (relative to training data). The state space is the collection of weights and the dynamical system used to tune a neural network is to follow along the gradient of Φ (using back-propagation or other algorithmic estimations) until reaching an equilibrium – a local minimum of Φ .



EXAMPLE: In Economics, equilibria are commonplace. Given simple supply and demand curves, there (usually) exists an equilibrium for price and quantity. Such an equilibrium can be viewed as an equilibrium for the (appropriate model of) market dynamics where sellers dynamically adjust price and quantity based on buyer response. In Game Theory, the notion of a Nash equilibrium is one where the state space is a space of probability distributions on strategies. A Nash equilibrium can be viewed as a critical point of the expected payoff function: a dynamical equilibrium of expected payoff.

EXAMPLE: If you have access to an old-fashioned scientific calculator with physical buttons, the following is a satisfying exercise. Input any random number, then hit the **cos** button, being careful to have set the calculator to “radians”. Hit that button again and again, and you will note a seeming convergence to one special number. Try it again with any initial condition, and the same number emerges. Is this a trick? No, it is an equilibrium, and it is **stable**.

STABILITY

Focusing on equilibria reveals a finer distinction. Most of the equilibria one “sees” are what one calls stable, meaning, roughly, that nearby initial conditions will converge to the equilibrium, or at least not depart from being near. In contrast, unstable equilibria have the property that a small perturbation to the equilibrium state leads to rush away from it. Balancing a long thin rod upright in the palm of your hand leads to an unstable equilibrium, as does carefully stacking one bowling ball atop another: it can be done in theory.

At the beginning of our story, stability is dichotomous, and equilibria are either stable or unstable. Later, as dimensions unfold, a richer theory of stability and instability ensues. As a foreshadowing, note that tipping back and balancing on two legs of a chair is unstable, but not nearly as unstable as trying to balance on one: instability comes in degrees.


BIFURCATIONS

Many dynamical systems arise as models, and models require choices. Often, certain details of the model are encoded as **parameters** – coefficients that do not

evolve with time, but which can be estimated or tweaked in advance of running the system. Examples include masses of bodies in a continuous-time model of gravitational dynamics, or the slopes of supply and demand curves in a discrete-time economic model. The set of “reasonable” parameters reside in a parameter space, usually $\mathbb{R}^p$, for some p , the number of parameters. One observes through small changes in the values of parameters that, although equilibria may change location, the number and types of equilibria remain constant. When there are qualitative changes in equilibria, they tend to occur at a specific location in state space and in parameter space. Such an event is called a **bifurcation**.

Near a bifurcation, a small change in a parameter may lead to a stable equilibrium becoming unstable, or the creation of new nearby equilibria, or the annihilation of a pair of nearby equilibria. One example is related to the way that a structural beam withstands more and more weight until, past a critical threshold, it buckles. Other examples include two-species population models, where a small change in the parameter related to species interaction can mean the difference between a stable, coexistence equilibrium and an unstable, competitive-exclusion equilibrium: peace vs. total war.

Remarkably, in the same way that equilibria can be classified, these bifurcations can themselves be categorized and identified. Some of the richest taxonomies of this text concern bifurcation types.



PERIODIC ORBITS


Equilibria, though important, are not the only dynamical features worth studying. Many dynamical systems exhibit behavior that is periodic. A **periodic orbit** is a solution to a dynamical system that repeats at a regular time interval $P > 0$: the state at time $t + P$ equals that at time t for all values of t . An equilibrium can be thought of as a periodic orbit with trivial period.

Periodic orbits abound in physical systems, especially those without friction: swinging pendula, vibrating strings, and orbiting satellites are periodic. Other systems can express periodicity in population sizes (resource oversupply, population expansion, overconsumption, resource depletion, population contraction, then underconsumption) or prices (high prices, suppliers flood the market, low prices, suppliers leave the market). Biological systems are especially prone to (roughly) periodic behaviors, due in part to circadian rhythms.

Periodic orbits, together with equilibria and bifurcations thereof, will form the skeleton on which qualitative dynamics hangs, in everything from the simplest to the most complex systems.

CHAOS

This text concerns *deterministic systems* – those for which the present determines the future without randomness or noise in how the dynamics unfolds. Although the future is set for a deterministic system, it is not necessarily known or knowable. Certain dynamical systems when simulated reveal complex unpredictable meanderings that cause nearby initial conditions to eventually disperse. This sensitive dependence on initial conditions is a hallmark of *chaos*.



Chaotic dynamics are brilliant & beautiful, and thus attract a swarm of clickbait butterflies, unattractive strangers, fractal financial scams, and pseudoclaims of founding a new science. The truth about chaos is simple: although individual paths through a chaotic system may seem unpredictable, the system itself may be completely analyzable.

EXAMPLE: It is best to begin demystification at once. Consider the discrete dynamical system which takes a number $0 \leq x < 1$ and multiplies it by 10, dropping the “whole

```
.314159265358979323...
.141592653589793238...
.415926535897932384...
.159265358979323846...
.592653589793238462...
.926535897932384626...
.265358979323846264...
.653589793238462643...
.535897932384626433...
```

number” part: $x_{n+1} = 10x_n \bmod 1$. Choose an initial condition x_0 between zero and one. If you represent this number as a decimal, then the dynamics consists of moving the decimal point one place to the right and dropping the first digit (now to the left of the decimal point). It is not an exaggeration to say that this is *the* example of a chaotic dynamical system.

Note how many fundamental features are present. There are several equilibria (9 in fact, corresponding to multiples of $1/9$). Repeating decimals (*e.g.*, $3/11$ or $1/7$) correspond to periodic orbits (of periods 2 and 6 respectively). Irrational numbers (which comprise “most” initial conditions) correspond to non-equilibrium, non-periodic orbits. If you begin with two initial conditions that are very close to one another (say, agreeing to within one billionth), then, after ten time steps it is impossible to say whether the orbits are close or what their future holds. This chaotic system loses information about state at every time step. All the hallmarks of chaotic dynamics are perfectly captured in this one example, an example we will meet again in Chapter 35.

CHAPTER 2 : CONTINUOUS & DISCRETE CALCULUS

DYNAMICAL SYSTEMS emerges from Mathematics and is best understood via mathematical tools. The first and greatest such toolset is Calculus, and, though its limits will quickly be reached, its language and principles are indispensable.

Just as there are two fundamental classes of dynamics – continuous and discrete time – there are two parallel calculi with which to solve problems. A review of each is in order.

CONTINUOUS CALCULUS

There are three fundamental concepts in calculus: asymptotics, derivatives, and integrals. Of these, the first is most misunderstood in the current curriculum and thus worthy of review.

Continuous single-variable calculus concerns functions of the form $f: \mathbb{R} \rightarrow \mathbb{R}$ which are sufficiently well-behaved. The first such restriction is that of continuity, meaning that for all $a \in \mathbb{R}$,

$$\lim_{x \rightarrow a} f(x) = f(a).$$

After continuity, differentiability a most desirable feature. Differentiability means that the derivative

$$f'(a) = \left. \frac{df}{dx} \right|_a = \lim_{h \rightarrow 0} \frac{f(a+h) - f(a)}{h}$$

exists at all a . The notation $f'(a)$, though standard, is suboptimal, as it does not disclose the variable with which the derivative is taken. This is to be avoided when, as is nearly always the case in dynamics, multiple variables are in play. In single-variable calculus, differentiation takes a function $f: \mathbb{R} \rightarrow \mathbb{R}$ to $f': \mathbb{R} \rightarrow \mathbb{R}$, which is especially convenient, as the process may be iterated.

To best deal with multiple (or “higher”) derivatives, it is convenient to use the notation of operators, where an **operator** is something like a “function of functions”. We denote by $D = d/dx$ the differentiation operator on (smooth) functions of one variable. The benefit of so doing lies in the ability to use powers: D^2 denotes the operator that differentiates twice; $D^0 = I$ is the **identity operator** that does nothing.

When using D , the variable should be clear, but at times may require discernment. Witness the general linear second order autonomous differential equation on $x(t)$,

$$a \frac{d^2 x}{dt^2} + b \frac{dx}{dt} + cx = 0 \quad \Rightarrow \quad (aD^2 + bD + cI)x = 0$$

Here, and for most of the rest of this text, the D operator represents differentiation with respect to t .

The differentiation operator is useful in much more than differential equations. In what follows, we assume functions to be **smooth**, meaning that all derivatives exist at all inputs, and, stronger still, **real-analytic**, meaning that Taylor series converge to the function.

The **Taylor expansion** of a function $f: \mathbb{R} \rightarrow \mathbb{R}$ about an input a is the series


$$f(x) = \sum_{k=0}^{\infty} \frac{1}{k!} D^k f|_a (x - a)^k$$

$$= f(a) + f'(a)(x - a) + \frac{1}{2!}f''(a)(x - a)^2 + \dots$$

where the differential operator is used to denote higher derivatives. It often helps to change coordinates using the variable $h = x - a$ to represent the distance from the expansion point. With this, the Taylor expansion is a polynomial series:

$$f(x + h) = \sum_{k=0}^{\infty} \frac{1}{k!} D^k f|_a h^k.$$

This will be *useful* to us, especially when we truncate to a polynomial of fixed size. The proper way to represent the leftover terms of a truncated series is to use **asymptotic notation**. For example, a few simple Taylor expansions of basic functions about $x = 0$ is given in truncated asymptotic form as follows:



$$e^x = \sum_{k=0}^{\infty} \frac{x^k}{k!} = 1 + x + \frac{x^2}{2} + \frac{x^3}{6} + O(x^4)$$

$$\sin x = \sum_{k=0}^{\infty} (-1)^k \frac{x^{2k+1}}{(2k+1)!} = x - \frac{x^3}{6} + \frac{x^5}{120} + O(x^7)$$

$$\cos x = \sum_{k=0}^{\infty} (-1)^k \frac{x^{2k}}{(2k)!} = 1 - \frac{x^2}{2} + \frac{x^4}{24} + O(x^6)$$

$$\frac{1}{1-x} = \sum_{k=0}^{\infty} x^k = 1 + x + x^2 + x^3 + x^4 + O(x^5) \quad : \quad |x| < 1$$

$$\ln(1+x) = \sum_{k=1}^{\infty} (-1)^{k+1} \frac{x^k}{k} = x - \frac{x^2}{2} + \frac{x^3}{3} + O(x^4) \quad : \quad |x| < 1$$

$$(1+x)^\alpha = \sum_{k=0}^{\infty} \binom{\alpha}{k} x^k = 1 + \alpha x + \frac{\alpha(\alpha-1)}{2} x^2 + O(x^3) \quad : \quad |x| < 1$$

For those who did not internalize these in calculus class, now is the time to do so. The big-O notation $O(x^k)$ means, informally, that all additional terms have degree k or greater – it is a trash can of size x^k . Formally, $f(x) \in O(g(x))$ as $x \rightarrow 0$ if and only if $|f(x)| \leq C|g(x)|$ for some constant C as $x \rightarrow 0$. This notation is uniquely well-suited both to learning and to applying calculus.

THE SHIFT OPERATOR

Besides the differentiation operator, D , and the identity operator, I , there is another operator of great importance in calculus and dynamics. The **shift operator**, E , shifts the input of a function f by one unit; that is,

$$Ef(x) = f(x + 1).$$

EXAMPLE: Powers of the shift operator allow for the following:

1. $E^2 f(x) = f(x + 2)$
2. $E^{-1} f(x) = f(x - 1)$
3. $E^h f(x) = f(x + h)$ for all real values of h

4. $E^a E^b = E^{a+b}$
5. $E^0 = I$

There is a deep connection between the shift operator E and the differentiation operator D . The following is a beautiful result: a candidate for one's favorite.

LEMMA: [The Exponential Lemma] $E = e^D$.

▷ Proof: Let f be real-analytic. Then, for all x ,

$$(e^D f)(x) = \sum_{k=0}^{\infty} \frac{D^k f|_x}{k!} = \sum_{k=0}^{\infty} \frac{D^k f|_x}{k!} 1^k = f(x+1) = (Ef)(x).$$

The Taylor formula justifies the penultimate equality and thus the conclusion. ◁

This proof may seem like a sleight. However, this humble lemma will repeatedly inform our observations about discrete versus continuous time dynamics: E is the canonical evolution operator in discrete-time dynamical systems.

DISCRETE CALCULUS

The discrete calculus is simpler, older, and much less well known outside of certain corners of numerical analysis and computational complexity theory. The only readily available references are either too old or idiosyncratic to be of much use to the typical student. Therefore, what follows is a summary of the theory that goes somewhat beyond what is directly relevant to the discrete-time dynamics of this text.

Discrete calculus is the calculus of functions which have a discrete input (typically $\mathbb{N}$ but sometimes $\mathbb{Z}$) and a continuous output. Such a function $a: \mathbb{N} \rightarrow \mathbb{R}$ is called a **sequence** and is usually denoted

$$a = (a_n) = (a_0, a_1, a_2, \dots).$$

Of the four core concepts of calculus, the first is least troublesome: there is no need to worry about limits, excepting the limit as the input goes to infinity, and this matches the definition in the continuous case of $f: \mathbb{R} \rightarrow \mathbb{R}$.

EXAMPLE: Consider the following sequences.

1. $(2n) = (0, 2, 4, 6, 8, \dots)$
2. $(n^2) = (0, 1, 4, 9, 16, \dots)$
3. $(\cos n) = (1, \cos 1, \cos 2, \cos 3, \dots)$
4. $(e^n) = (1, e, e^2, e^3, \dots)$

Which have graphs that allow you to recognize the smooth analogue?

Derivatives are less obvious, as the usual $\epsilon - \delta$ definition will not do. Thus, there are two types of discrete derivative, called the **forward** and **backward difference**:

$$\Delta a = (a_{n+1} - a_n) \quad : \quad \nabla a = (a_n - a_{n-1}).$$

These of course fit the intuition of a derivative as a change in output per change in input, as well as giving the slope of a “tangent” to the graph, to the right/left respectively. We will use the forward difference exclusively.

EXAMPLE: Consider the forward differences of the following sequences.

1. $\Delta(n^2) = ((n+1)^2 - n^2) = (2n+1)$
2. $\Delta(n^3) = ((n+1)^3 - n^3) = (3n^2 + 3n + 1)$

3. $\Delta(e^n) = (e^{n+1} - e^n) = (e^n(e - 1))$
4. $\Delta(2^n) = (2^{n+1} - 2^n) = (2^n)$

In what way do these resemble continuous derivatives? How are they different?

With a derivative comes the ability to write (and solve) differential equations. The simplest are linear differential equations of the form $\Delta x = \lambda x$ where λ is a constant. These have as solutions $(x_n) = ((\lambda - 1)\lambda^n)$. Note that such a sequence equals its difference if and only if $\lambda = 2$. That is, we have

$$\Delta x = x \Rightarrow x = (2^n).$$

This suggests that in discrete calculus “ $e = 2$ ” and that the discrete exponential function is 2^n .

The difference operator can be iterated to yield the second derivative Δ^2 and so on. This is best done through the language of operators. One notes that:

$$\Delta = E - I.$$

EXAMPLE: Higher derivatives are easily computed via algebra.

$$\Delta^3 x = (E - I)^3 x = (E^3 - 3E^2 + 3E + I)x = (x_{n+3} - 3x_{n+2} + 3x_{n+1} - x_n).$$

Integrals can be presented as a formal inverse to derivatives (an indefinite integral), using the geometric series in a somewhat suspicious manner:

$$\Delta^{-1} = (E - I)^{-1} = -(I - E)^{-1} = -\sum_{k=0}^{\infty} E^k.$$

It is much safer to discuss the definite integral, which is simply a finite sum. The Fundamental Theorem of Integral Calculus then has a discrete version which is more familiar as a telescoping sum:

$$\sum_{k=a}^b (\Delta x)_k = x_{b+1} - x_a.$$

FALLING POWERS

There is a notational maneuver for making sense of discrete calculus: this is the *falling power*. For $k > 0 \in \mathbb{N}$, the expression $x^{\underline{k}}$, that is “ x to the falling k ” is

$$x^{\underline{k}} = x(x - 1)(x - 2) \cdots (x - k + 1).$$

For an integer n , this expression simplifies to $n^{\underline{k}} = n!/(n - k)!$ and in either case is a degree k polynomial. It is wise to set $x^0 = 1$, and we, being wise, shall do so.

EXAMPLE: Consider the following discrete monomials using falling powers:

1. $(n^{\underline{2}}) = (0^{\underline{2}}, 1^{\underline{2}}, 2^{\underline{2}}, \dots, n^{\underline{2}}, \dots) = (0, 0, 2, 6, 12, 20, \dots, n(n - 1), \dots)$
2. $(n^{\underline{3}}) = (0^{\underline{3}}, 1^{\underline{3}}, 2^{\underline{3}}, \dots, n^{\underline{3}}, \dots) = (0, 0, 0, 6, 24, 60, \dots, n(n - 1)(n - 2), \dots)$
3. $(n^{\underline{k}}) = (0^{\underline{k}}, 1^{\underline{k}}, 2^{\underline{k}}, \dots, n^{\underline{k}}, \dots)$

The notation is built to make discrete calculus work.

LEMMA: $\Delta(n^{\underline{k}}) = (k n^{\underline{k-1}})$.

▷ Proof: One computes:

$$\Delta(n^{\underline{k}}) = ((n + 1)^{\underline{k}} - n^{\underline{k}}) = \left(\frac{(n + 1)!}{(n + 1 - k)!} - \frac{n!}{(n - k)!} \right)$$

$$= \left(\frac{((n+1) - (n+1-k))n!}{(n+1-k)!} \right) = \left(\frac{k n!}{(n-(k-1))!} \right) = (k n^{\underline{k-1}}).$$

The cumbersome simplification uses the common denominator $(n+1-k)!$. $\triangleleft$
This makes it relatively simple to compute discrete derivatives and antiderivatives of polynomial functions using falling powers.

DISCRETE TAYLOR EXPANSION

Taylor series are one of the nicer applications of falling powers. Taylor expansion (& much of calculus) evolves out of the classical **Binomial Theorem**, which states:

$$(a+b)^n = \sum_{k=0}^n \binom{n}{k} a^{n-k} b^k = \sum_{k=0}^n \frac{n!}{(n-k)! k!} a^{n-k} b^k.$$

The discrete calculus version of Taylor expansion uses this to evaluate powers of the shift operator. Given a sequence, x , if we want to know the n^{th} term, x_n , using only information about x_0 and the “derivatives” of x at x_0 , then, using the fact that $\Delta = E - I$, we can write:

$$x_n = (E^n x)_0 = (I + \Delta)^n x|_0 = \sum_{k=0}^n \binom{n}{k} I^{n-k} \Delta^k x|_0 = \sum_{k=0}^n \frac{\Delta^k x|_0}{k!} n^{\underline{k}}.$$

This is a finite, discrete-calculus version of Taylor expansion about zero. The result (using different notation) is due to Newton and predates Taylor series.

Neither this, nor falling powers in general will be needed in the remainder of this text. However, deep familiarity with the relationship between the difference operator, Δ , the shift, E , and the differentiation operator, D , is crucial in dynamical systems.

CHAPTER 3 : EQUILIBRIA & STABILITY

THE PRIMAL dynamical features in any system are its equilibria, or constant solutions. Our first task after introducing dynamical systems is to learn how to find and classify its equilibria. Many of the major themes of the text – nonlinearity, linearization, and classification – are introduced here.

LINEAR & NONLINEAR DYNAMICS

We work with autonomous 1-D dynamical systems of the following form:

1. CONTINUOUS TIME : $Dx = f(x)$ for $x = x(t)$ a continuous function.
2. DISCRETE TIME : $Ex = f(x)$ for $x = (x_n)$ a discrete function.

The simplest such systems are linear dynamics: $f(x) = \lambda x$ for λ a constant. These can be solved explicitly:

$$Dx = \lambda x \Rightarrow x(t) = e^{\lambda t} x_0 \quad : \quad Ex = \lambda x \Rightarrow x_n = \lambda^n x_0.$$


This will soon be useful. It is possible though nontrivial to solve certain nonlinear dynamical systems explicitly using tools from calculus.

EXAMPLE: A classic example of a continuous-time nonlinear system solved in Calculus classes is the logistic model.

$$\frac{dx}{dt} = rx(K - x)$$

Here, $r, K > 0$ are constants representing a reproduction rate, r , and a carrying capacity, K . Using a combination of separation, integration by partial fractions, and no small amount of algebra, one shows that

$$x(t) = \frac{Kx_0}{x_0 + (K - x_0)e^{-rKt}}$$



One verifies that at $t = 0$, the population size matches the initial condition x_0 ; furthermore, for $x_0 > 0$ and $t \rightarrow \infty$, the population tends to $x(t) \rightarrow K$.

EXAMPLE: A classic example of a discrete-time (slightly) nonlinear system comes from the Towers of Hanoi problem, which enumerates how many moves x_n it takes to move a stack of n ordered discs from one peg to another using moves that preserve the order of the stack elements. The recurrence relation (obtained by induction) is

$$x_{n+1} = 2x_n + 1.$$


This is nearly a linear model: it is more properly called *affine* (linear-plus-constant). One can solve this by “guessing” a solution based on a few iterations. Better is to solve by a

discrete form of integration and “ u -substitution”. Let $u_n = x_{n+1}$. Then, converting to u gives: $Eu = Ex + 1 = 2x + 2 = 2(x + 1) = 2u$.

This u -substitution converts the affine dynamics in x to linear dynamics in u . The solution is therefore:

$$\begin{aligned} u_n = 2^n u_0 &\Rightarrow x_{n+1} = 2^n(x_0 + 1) \\ &\Rightarrow x_n = 2^n x_0 + 2^n - 1. \end{aligned}$$

With the traditional interpretation of the Towers of Hanoi problem, $x_0 = 0$, yielding $x_n = 2^n - 1$.



The above examples are the exception: nonlinear dynamical systems are simply too difficult to solve in general. Consider the continuous-time system $Dx = f(x)$. The general solution is given by separation, integration, and algebra as follows:

$$\frac{dx}{dt} = f(x) \Rightarrow \frac{dx}{f(x)} = dt \Rightarrow \int \frac{dx}{f(x)} = t + C \Rightarrow x(t) = \sqrt[3]{(t+C)^2}$$

In all but a few toy cases, it is either impossible to integrate the reciprocal of f , or impossible to invert this result for x as a function of t . The discrete time case is no less difficult for f nonlinear. A different strategy is needed.

EQUILIBRIA & STABILITY

Equilibria are special solutions to dynamical systems. The definition is unchanging between discrete and continuous time: an equilibrium is a constant solution:

- ▷ CONTINUOUS TIME : $Da = f(a) = 0$ for a a constant.
- ▷ DISCRETE TIME : $Ea = f(a) = a$ for a a constant.

Examples of equilibria abound, as hinted in Chapter 1. We will focus on equilibria of specific 1-D models in continuous and discrete time; but the reader should think more broadly about what equilibria mean in contexts ranging from economic to social, psychological, and more.

Recall that equilibria can come in several types, roughly classified in 1-D by a notion of stability. For the remainder of this Volume, we declare a **stable equilibrium** to be one which attracts all nearby initial conditions; an **unstable equilibrium** repels all nearby initial conditions – it is stable in reverse-time. An equilibrium which is neither stable nor unstable will be called **degenerate**. This trichotomy is simplistic, and we will need to update our definitions once we leave the safety of 1-D; but for now, stable and unstable equilibria will be identified with attracting and repelling behaviors.

LINEARIZATION & CLASSIFICATION

Nonlinear dynamics are difficult if not impossible to solve explicitly. Our strategy will be to identify the equilibria and classify them as stable or unstable (if possible). This will be accomplished using linearization to convert the nonlinear global system to a linearized local system. This is a fantastic generalization of the concept of Taylor expansion; we will not be linearizing a function so much as linearizing a dynamical system.

CONTINUOUS TIME : Let a be an equilibrium, and let x be near a , with $h = x - a$ an indication of how close to the equilibrium one is. Assume $x = x(t)$

satisfies the differential equation $Dx = f(x)$. Then the distance to a , $h(t)$, evolves as

$$\frac{dh}{dt} = \frac{d}{dt}(x - a) = \frac{dx}{dt} - 0 = f(x) = f(a + h) = f(a) + \left. \frac{df}{dx} \right|_a h + O(h^2),$$

thanks to Taylor expansion of f about a . Since a is an equilibrium, this means that

$$Dh \approx \left. \frac{df}{dx} \right|_a h.$$

This is the linearization of $Dx = f(x)$ about the equilibrium $x = a$, and it is a linear dynamical system with solution

$$h(t) = e^{\lambda t} h(0) : \lambda = \left. \frac{df}{dx} \right|_a$$

This means that perturbations to the equilibrium grow exponentially if the derivative df/dx at $x = a$ is positive and decay exponentially if this derivative is negative. Caveat: there are multiple derivatives at play when linearizing. This is the derivative with respect to state, x , not time, t . Avoiding the use of “prime” notation for derivatives is highly encouraged.

DISCRETE TIME : Let a be an equilibrium of $Ex = f(x)$, and let x be near a , with $h = x - a$ an indication of how close to the equilibrium one is. Assume $x = (x_n)$ satisfies the recurrence $Ex = f(x)$. Then the distance to a , $h = (h_n)$, satisfies

$$E(a + h) = Ea + Eh = a + Eh = f(x) = f(a + h) = f(a) + \left. \frac{df}{dx} \right|_a h + O(h^2),$$

thanks to Taylor expansion of f about a . Since a is an equilibrium, $f(a) = a$ and

$$a + Eh \approx a + \left. \frac{df}{dx} \right|_a h \Rightarrow Eh \approx \left. \frac{df}{dx} \right|_a h.$$

This is the linearization of $Ex = f(x)$ about the equilibrium $x = a$. The solution to the linearized system is

$$h_n = \lambda^n h_0 : \lambda = \left. \frac{df}{dx} \right|_a$$

Note the difference: perturbations to the equilibrium grow exponentially if the derivative df/dx at $x = a$ is greater than one in absolute value. The perturbations decay if this derivative is less than one in absolute value.

THE STABILITY CRITERION

By linearizing about an equilibrium and using what we know about explicit solutions to linear systems, we have proved the first major result of this subject: a criterion for determining the stability of an equilibrium.

STABILITY CRITERION, CONTINUOUS TIME

Let a be an equilibrium of the dynamical system $Dx = f(x)$, and denote by λ the derivative $\lambda = df/dx|_a$ evaluated at the equilibrium. Then a is

- ▷ STABLE if $\lambda < 0$
- ▷ UNSTABLE if $\lambda > 0$

▷ DEGENERATE if $\lambda = 0$

STABILITY CRITERION, DISCRETE TIME

Let a be an equilibrium of the dynamical system $Ex = f(x)$, and denote by λ the derivative $\lambda = df/dx|_a$ at the equilibrium. Then a is

▷ STABLE if $|\lambda| < 1$

▷ UNSTABLE if $|\lambda| > 1$

▷ DEGENERATE if $|\lambda| = 1$

It is worth contemplating the similarities and differences between these stability criteria. Both use the coefficient, λ , of the first-order term in the Taylor expansion of f , but the discrete-time criterion seems more complex, with an absolute value thrown in.

The best answer to why these criteria differ in this precise manner comes from understanding the relationship between the continuous-time evolution operator, D , and the discrete-time evolution operator, E . Recall from Chapter 2 that $E = e^D$. What happens when you exponentiate the stability criterion for D ? When $\lambda > 0$, then $e^\lambda > 1$, and when $\lambda < 0$, then $0 < e^\lambda < 1$. The addition of the absolute value in the discrete-time case is an as-yet-unexplained complexity that will make more sense when we move up a dimension.

CLASSIFYING EQUILIBRIA

Let us put our skills to work and consider a few simple examples.

LOGISTIC MODELS

Recall the logistic model for populations in continuous time:

$$\frac{dx}{dt} = f(x) = rx(K - x) \quad : \quad r, K > 0.$$

This has equilibria at $x = 0$ and at $x = K$. Stability is determined by the derivative

$$\frac{df}{dx} = rK - 2rx.$$

At $x = 0$ this evaluates to $rK > 0$: an unstable equilibrium. At $x = K$, we have derivative $-rK < 0$: a stable equilibrium. In this simple model, the linear term rKx accounts for the exponential growth of small populations. The quadratic term $-rx^2$ retards growth when the population size is sufficiently large. The result is a stable population size, K , to which any nonzero initial condition converges.

A discrete time version of the logistic model appears similar:

$$x_{n+1} = f(x_n) = x_n + rx_n(K - x_n) \quad : \quad r, C > 0.$$

This also has equilibria at $x = 0, K$; however, there are some subtle differences. Since the derivative of the right-hand side is

$$\frac{df}{dx} = 1 + rK - 2rx.$$

Then, as before, the equilibrium at $x = 0$ is unstable for all $r, K > 0$. At $x = K$ the derivative evaluates to $1 - rK$, which is a stable equilibrium so long as $rK < 2$.

BUCKLING BEAM

Recall from Chapter 1 the phenomenon of the buckling beam. The deflection, $x(t)$, at the midpoint of the beam is given by the differential equation

$$\frac{dx}{dt} = f(x) = \lambda x - x^3,$$


where $\lambda > 0$ is the force at the end. The equilibria occur where $x(\lambda - x^2) = 0$. That is, there are three equilibria, at $x = 0, \pm\sqrt{\lambda}$. The stabilities of these equilibria are determined by the derivative of the right hand side with respect to x ,

$$\frac{df}{dx} = \lambda - 3x^2.$$

At $x = 0$, this evaluates to $\lambda > 0$, and the equilibrium is unstable. At $x = \pm\sqrt{\lambda}$, the derivative evaluates to $-2\lambda < 0$; these equilibria are stable.

COSINE BUTTON-SMASH

Recall from Chapter 1 that inputting a random number into a calculator and pressing the cosine button repeatedly yields a convergence to a mysterious constant $a \approx 0.73908 \dots$. Is this some heretofore hidden universal constant? No, it is simply a stable equilibrium of the discrete-time system $Ex = f(x) = \cos x$. Solving the equation $\cos x = x$ (using, say, Newton's method) yields this a as the unique equilibrium. The derivative $df/dx = \sin x$ evaluated at $x = a$ is approximately -0.67361 . Since the absolute value is less than one, this equilibrium is stable.



NEWTON's METHOD

We have just invoked the use of Newton's method: it is work dwelling on this. Given a differentiable function $g(x)$, one can use Newton's method to find a root. Recall that Newton's method involves choosing an initial condition, x_0 , and iterating using the formula

$$x_{n+1} = x_n - \frac{g(x_n)}{g'(x_n)},$$

where, for consistency with basic calculus texts, g' is used instead of the more precise expression dg/dx . You should now recognize this as a discrete time dynamical system with right hand side $f(x) = x - g/g'$. This has an equilibrium where $x = x - g(x)/g'(x)$; that is, $x = a$, where $g(a) = 0$ and $g'(a) \neq 0$. Is such a root a stable equilibrium? Quite. One computes

$$\left. \frac{df}{dx} \right|_a = 1 - \frac{g'(a)g'(a) - g(a)g''(a)}{(g'(a))^2} = 1 - \left(\frac{g'(a)}{g'(a)} \right)^2 = 0.$$

This explains why Newton's method works so very well in a neighborhood of a root. The equilibrium is what one calls **superstable**.

THE MEANS OF GOLD & SILVER

Consider the following continued fraction and call the resulting limit Φ .

$$\Phi = \frac{1}{1 + \frac{1}{1 + \frac{1}{1 + \frac{1}{1 + \frac{1}{1 + \frac{1}{1 + \dots}}}}}}$$

This is, indeed the famous **golden mean**, equal to $(1 + \sqrt{5})/2$. We can represent this continued fraction in terms of a discrete time dynamical system as follows.

$$x_{n+1} = \frac{1}{1 + x_n} \quad : \quad x_0 = 0.$$

This dynamical system has equilibria where $a(1 + a) = 1$; that is, via the quadratic formula, at $(1 \pm \sqrt{5})/2$. These equilibria are (up to a sign) the golden and silver means, respectively. What of their stability? One computes:

$$\left. \frac{d}{dx} \left(\frac{1}{1 + x} \right) \right|_a = -(1 + a)^{-2} = -\frac{4}{3 \pm \sqrt{5}}.$$

At the golden mean $(1 + \sqrt{5})/2$, the derivative is between -1 and 0 : a stable equilibrium. In contrast, at the silver mean $(1 - \sqrt{5})/2$, the derivative is less than -1 , and this is an unstable equilibrium. The continued fraction converges to the stable equilibrium Φ .

CHAPTER 4 : GRAPHICAL REPRESENTATIONS


THUS FAR our approach to dynamical systems has been more analytical and quantitative than qualitative or geometric. This is appropriate at the beginning, to reinforce proper technique. However, as we progress through the subject, turning to more sophisticated and realistic models, multiple modes of thinking and processing of dynamics will be invaluable. This is especially true when it comes to visual representations of dynamics.

There are several approaches to visualizing dynamical systems, almost all of which are better served by dynamic than by static imagery. While the videos that pair with this text should be viewed as a primary reference, the flat imagery here printed can serve as a proxy for the richer content.

As ever, our discussion splits between continuous and discrete time systems, with the reader being encouraged to contemplate the passage from one to the other.


CONTINUOUS TIME: TEMPORAL PLOTS

The simplest type of plot associated to a continuous-time dynamical system is that which is popular in basic calculus classes. Given $dx/dt = f(x)$, one plots a solution $x(t)$ as a function of t , with, perhaps, multiple initial conditions. Given such a plot, one can identify stable or unstable equilibria easily as straight line solutions. Stability is indicated as a type of contraction or expansion of nearby curves along the time axis.



This is fine. However, one cannot but wonder how such a figure would generalize, either to higher dimensional state spaces or to systems with parameters (*cf.* Chapter 7). It seems a prodigal use of space to plot the particulars of every solution.

CONTINUOUS TIME: DIAGRAMS



A more parsimonious representation comes from suppressing the time variable and instead plotting the derivative, $f(x)$, against x . The terminology of such figures is muddled in the literature. For simplicity, we will call such a plot a **diagram**. (Though not a terribly descriptive name, it is modest, and modesty is a virtue.)

Diagrams are informative. First, one sees the equilibria as the zero-set where the graph of f crosses the horizontal axis (where $dx/dt = 0$). Next, stability and the stability criterion are easily visualized. For a

an equilibrium, linearizing f about a is represented by the tangent line to the graph of f at a . A tangent line with negative slope implies that $f(a - \epsilon) > 0$


and $f(a + \epsilon) < 0$ for small $\epsilon > 0$: this is a stable equilibrium, since the sign of $f(x)$ determines whether x is increasing or decreasing over time.

One sees well from such a diagram that, so long as f is continuous, the behavior of f away from the zero-set is irrelevant to the qualitative dynamics: what matters is the sign of dx/dt , indicated by the graph of f . The behavior of the system hangs on a skeleton of equilibria.

This is worth repetition. A diagram makes clear that what matters in continuous time is the equilibria and the local behavior near equilibria. Everything else influence only the speed of solutions and not the qualitative features.

CONTINUOUS TIME: VECTOR FIELDS & FLOWS

One can take the suppression of the time variable to an extreme limit. Consider the equation $Dx = f(x)$ as defining a **vector field** on $\mathbb{R}$. Students will recall from vector calculus the ubiquity & utility of vector fields in 2-D and 3-D. This is a 1-D vector field, illustrated by drawing a small arrow (or perhaps just an arrowhead) at each point, with equilibria prominently displayed with a dot.



at each point, with equilibria prominently displayed with a dot.


The lesson, again, is that the directionality of the arrows

matters much more than the magnitude. In later volumes of this text, the vector field approach will rise to a place of prominence as a visualization tool.

DISCRETE TIME: DIAGRAMS


Discrete time systems require an entirely novel method of visualization. One of the most popular is popularly called a *cobweb plot*. This terminology, as is common in dynamical systems, perfectly encapsulates how not to name a thing. We will instead use the plain and simple [discrete time] **diagram** to denote the following.

Given a system of the form $x_{n+1} = f(x_n)$, plot the graph of f on a plane with axes x_n and x_{n+1} . Each input-output pair (x_n, x_{n+1}) becomes a point on the graph, and an orbit of the system is a sequence of such points. To show the progression from the pair $(x_{n-1}, x_n) \rightarrow (x_n, x_{n+1})$, draw the pair of (directed) axis-aligned segments



$$(x_{n-1}, x_n) \rightarrow (x_n, x_n) \rightarrow (x_n, x_{n+1}).$$

This gives a diagram that shows the progression of discrete time. Equilibria are indicated by places where the graph of f intersects the diagonal where $x_{n+1} = x_n$. The stability criterion also becomes clear with the proper diagram. At an equilibrium a , if the derivative $|df/dx| < 1$, then one sees from the diagram



that nearby initial conditions are drawn in to the equilibrium point, whereas they are repelled for slope larger than one (in absolute value).

One notes that, unlike in the case of continuous time, discrete-time diagrams can have very different-looking orbits for

nearby initial conditions. The way that the stair-step graphs weave can be confusing: it is best (as with many things in dynamical systems) to watch it unfold in video.

EXAMPLE: In a discrete-time logistic model of the form

$$x_{n+1} = \frac{5}{2}x_n(1 - x_n),$$

there are equilibria at $x = 0$ and $x = 3/5$. The derivative of the right hand side evaluates at $x = 0$ to $5/2$ and at $x = 3/5$ to $-1/2$. Thus, the equilibrium at zero is unstable, and the equilibrium at $3/5$ is stable. One can see this in the diagram, localized about each equilibrium. At the equilibrium with negative derivative, it is as if the stair-step graph is spiraling into the equilibrium.


One sees as well that there is a substantial difference between positive and negative derivatives. The latter implies that orbits of nearby points *flip* about the equilibrium, not unlike the manner in which an alternating series converges. This will be an important observation in time.

DIAGRAMS OF CHAOS

Our use of visualization for dynamics illuminates one distinction between systems in continuous versus discrete time. In continuous time systems, the story is complete: given any reasonable function $f(x)$, a brief scan of the graph of f gives a clear picture of the dynamics. This is not so in discrete time.

Examining diagrams near stable or unstable equilibria is illuminating and clear. However, many examples of systems of the form $Ex = f(x)$ have diagrams which seem complex and bizarre. Some initial conditions appear to settle into repeated patterns; others appear to wander aimlessly, creating a tangled mess. Such is neither an accident of the representation nor an exceptional event.

Discrete-time dynamics in 1-D implicates much more than equilibria: periodic and chaotic dynamics abound in such systems. In the next chapter, we turn to periodic phenomena. Explaining chaotic dynamics in full is delayed until Volume 4; in Chapter 9, a brief foreshadowing will appear.



CHAPTER 5 : PERIODIC ORBITS

EQUILIBRIA are not the end. There are many more dynamical phenomena of interest than stationary points. This chapter gives a brief introduction to periodic phenomena in dynamics. The world runs in circles, and all living things express periodic behaviors, from sleep cycles and heartbeats to locomotive gaits.

PERIODIC ORBITS

A *periodic orbit* of a dynamical system is a solution that comes back to initial condition after a certain (positive) time. Specifically,

- ▷ CONTINUOUS TIME : $x(t + P) = x(t)$ for all t and some fixed $P > 0$.
- ▷ DISCRETE TIME : $x_{n+P} = x_n$ for all n and some fixed $P > 0$.

In each case, the minimal such P is denoted the *period* of the orbit. One can view an equilibrium as a trivial periodic orbit of period zero (in continuous-time; in discrete-time, period one). One usually reserves the term *periodic orbit* for the non-equilibrium case.

CONTINUOUS TIME

There are no nontrivial periodic orbits in continuous-time autonomous dynamics on $\mathbb{R}^1$. This seems intuitively obvious – how can you come back to where you start if you follow instructions on where to go based on location alone? However, writing out a proof of this can be a bit unintuitive, in part because the result is topological in nature, relying crucially on continuity properties.

LEMMA: In any system of the form $Dx = f(x)$, where f is continuous on $\mathbb{R}^1$, there are no periodic orbits.

▷ Proof: Assume that $x(t)$ is a solution to $Dx = f(x)$ satisfying $x(t + P) = x(t)$ for all t and some minimal $P > 0$. Then $f(x(t))$ never vanishes (else x would be an equilibrium). However, by the Fundamental Theorem of Integral Calculus, continuity of f , and periodicity of x ,

$$0 \neq \int_{t=0}^P f(x(t))dt = \int_{t=0}^P \frac{dx}{dt} dt = x(P) - x(0) = 0.$$

This is a contradiction. ◁

This is a significant limitation of 1-D continuous-time systems, since, as observed in Chapter 1, so much of the observable (and especially biological) world mirrors periodic dynamics. The way forward to more physically realistic settings is to add dimensions, as we shall do in Volumes 2-3 of this text.

DISCRETE TIME

In discrete time, we can recover a nonexistence result under the right assumption:

LEMMA: In any system of the form $Ex = f(x)$ on $\mathbb{R}^1$ with f increasing, there are no periodic orbits.

▷ Proof: Assume that such exists: $x(n + P) = x(n)$ for all n . Since $P > 1$, either $x_1 > x_0$ or $x_1 < x_0$. In either case, monotonicity of f implies that the sequence (x_n) is likewise monotone: contradiction. For a proof using the discrete-calculus version of the FTIC, consider this elegant analogue of the continuous-time proof:


$$0 \neq \sum_{n=0}^{P-1} (f(x_n) - x_n) = \sum_{n=0}^{P-1} (\Delta x)_n = x(P) - x(0) = 0,$$

where the first inequality comes from monotonicity of f . $\triangleleft$

FINDING PERIODIC ORBITS

Discrete time systems of the form $Ex = f(x)$ where f is non-increasing or non-injective can and do possess periodic orbits (& much more). Consider what happens in the following examples.

- ▷ If $f(x) = -x^3$, then the points $\{\pm 1\}$ comprise a periodic orbit of period two. This f is injective, but nonincreasing.
- ▷ If $f(x) = \frac{7}{2}x(1-x)$, then the points $\{\frac{3}{7}, \frac{6}{7}\}$ comprise a periodic orbit of period two.
- ▷ If $f(x) = 3.831874x(1-x)$, then there is a period-3 orbit with periodic points at (approximately) $\{0.5, 0.957969, 0.1542898\}$. With more decimal places, one can describe an exact period-3 orbit: try iterating these numbers on a calculator to see if you can get more digits of accuracy.




Notice that in discrete-time, a period- P orbit consists of P points, with the dynamics cycling through them in a set order that is not necessarily monotone. It is important to distinguish between the periodic points (of which there are P) and the periodic orbit (a single entity).

How one finds periodic orbits is perhaps not obvious. One way – that presages a deep idea in dynamical systems – stems from the observation that each point x_n in a period- P orbit of $Ex = f(x)$ is an equilibrium of the composition

$$f^{(P)} = f \circ f \circ \dots \circ f$$

of f with itself P times. Why? The language of operators again assists. By definition, x is a period- P orbit if,

$$\begin{aligned} x_n = x_{n+P} &\Rightarrow x = E^P x \Rightarrow x_n \\ &= E(E(\dots(Ex_n)\dots)) \\ &= f^{(P)}(x_n) \end{aligned}$$




It is important to note that not all equilibria of $f^{(P)}$ are period- P orbits – any other orbits of period dividing P (including 1) will also show up as equilibria.

One advantage of this method comes from trying to define stability of periodic orbits. Periodic orbits can be stable or unstable, based on whether initial conditions close to one of the points are attracted to the orbit or repelled away. This is inherited from the stability of one (and thus all) of the corresponding equilibria of $f^{(P)}$. In the above examples of periodic orbits, which are stable?

A PERFECT CIRCLE

There is one way to circumvent the absence of periodic orbits in 1-D continuous-time: modify the domain from the reals $\mathbb{R}$ to a circle. The mathematician's circle is a topological space denoted $\mathbb{S}^1$. This can be defined in several equivalent ways:

- ▷ $\mathbb{S}^1 = \{(x, y) \in \mathbb{R}^2 : x^2 + y^2 = 1\}$
- ▷ $\mathbb{S}^1 = \{z \in \mathbb{C} : |z| = 1\}$
- ▷ $\mathbb{S}^1 = \{e^{i\theta} \in \mathbb{C} : \theta \in \mathbb{R}\}$
- ▷ $\mathbb{S}^1 = [0, 2\pi] / \{0 = 2\pi\}$
- ▷ $\mathbb{S}^1 = \mathbb{R}/\mathbb{Z}$



The first three identify the circle as the unit circle in the real plane or complex line, and these are obviously equivalent (thanks to Euler's formula). The latter two definitions are as **identification** or **quotient spaces**, where one identifies certain points. One simple way to obtain a circle is to start with the closed interval $[0, 2\pi]$ and “glue” the two endpoints together by declaring that $0 = 2\pi$. In geometry, trigonometry, or calculus, one often works on the circle in this model, using angles modulo 2π . A linear rescaling from mod 2π to mod 1 given the reals modulo the integers, denoted $\mathbb{R}/\mathbb{Z}$. This presentation means that two real numbers are declared equivalent if they differ by an integer. Such equivalence classes parametrize a perfect circle.

Which of these is the *ideal* definition of $\mathbb{S}^1$? Which should be used in practice? That is a matter of choice. What is true is that each of the above descriptions of the circle are *topologically equivalent*, a concept of significant importance in its own right. The following definition will be of use on multiple occasions.

DEFINITION: A **homeomorphism** $f: X \rightarrow Y$ between spaces is a continuous function with continuous inverse.

Unfortunately, there is insufficient room to unfold all terms (*space* and *continuous* are doing the heavy work here). It is a challenging but ultimately helpful exercise to internalize what homeomorphism entails and why the different descriptions of $\mathbb{S}^1$ are all homeomorphic.

In practice, using an angular variable θ is a convenient way to coordinatize $\mathbb{S}^1$. The continuous-time dynamical system

$$\frac{d\theta}{dt} = \omega \quad \Leftrightarrow \quad D\theta = \omega,$$

for $\omega \neq 0$ a constant, is easily solved to yield $\theta(t) = \omega t + \theta_0$. This is a periodic orbit, since θ is defined modulo 1 (or 2π , depending on preference of coordinates). This is one simple way to get a periodic orbit in 1-D continuous-time. The general situation in continuous-time is similar:

LEMMA: Given $D\theta = f(\theta)$, for $f: \mathbb{S}^1 \rightarrow \mathbb{R}$, either (1) there are equilibria and no periodic orbits; or (2) there are no equilibria, and every orbit is periodic.

▷ *Idea of Proof:* To show (1), choose an equilibrium point. It cannot participate in a periodic orbit. Then, “puncture” $\mathbb{S}^1$ by removing the equilibrium, resulting in a dynamical system on $\mathbb{R}$ that we know has no periodic orbits. Showing (2) requires knowing that a circle is **compact**, so that f attains both global


minimum and global maximum, each of the same sign. Moving along the circle at that speed (or greater) brings one inexorably back to any initial condition. $\triangleleft$
 The situation is different for discrete-time dynamics on $\mathbb{S}^1$.

ROTATION MAPS

Discrete-time dynamics on a circle form a fascinating set of examples. Consider the simple rotation on $\mathbb{S}^1$ in coordinates $\theta \in \mathbb{R}/\mathbb{Z}$:

$$\theta_{n+1} = \theta_n + \omega \quad \Leftrightarrow \quad E\theta = \theta + \omega,$$

for $\omega \neq 0$ a constant. This has solution $\theta_n = \theta_0 + n\omega \bmod 1$. Everything depends on the spin constant ω . For $\omega \in \mathbb{Q}$ rational, all orbits are periodic, since multiplication of a rational (p/q) by a sufficiently large natural number (q) yields an integer. However, if $\omega \notin \mathbb{Q}$ is irrational, then, by the same argument, there are no periodic orbits whatsoever. In fact, each orbit in such an **irrational rotation** is dense within $\mathbb{S}^1$, wandering forever within the interstices of an unrepeating decimal.



This is a foreshadowing of what can occur when encountering a nonlinear system on a circle. Our exploration of rotations, rational and irrational, should lead to several questions about nonlinear systems of the form $E\theta = f(\theta)$.

- I. Is it possible for such a system to have multiple periodic orbits of different periods?
2. Is it possible to have both periodic orbits and dense orbits?
3. If you do not have periodic orbits or equilibria, must the orbits be dense?
4. Do the answers to the above change depending on invertibility of f ?

TOPOLOGICAL CONJUGACY

One of the most important, subtle, and pernicious issues is what it means for two mathematical entities to be *the same*. Mathematicians have a variety of tools to deal with issues of equality and equivalence. In the setting of dynamical systems, there is an equivalence relation which is acknowledged as the proper comparison. This is called **topological conjugacy**.

Given a pair of (discrete-time) dynamical systems $Ex = f(x)$ and $Ey = g(y)$ on spaces X and Y respectively, one views them as equivalent if there is a topological conjugacy between them: a homeomorphism $\varphi: X \rightarrow Y$ such that $\varphi \circ f = g \circ \varphi$. This condition means that φ takes orbits on X to orbits on Y bijectively: the dynamics match up. For continuous-time systems, it suffices to define a topological conjugacy for $e^{\epsilon D}$, the “flow forward by time ϵ ” map.

A number of results on circle maps require this notion of conjugacy to state. This is one of the most famous and subtle such results:

DENJOY’S THEOREM: If $E\theta = f(\theta)$, for $f: \mathbb{S}^1 \rightarrow \mathbb{S}^1$ a twice-differentiable function, and if this system has neither equilibria nor periodic orbits, then it is topologically conjugate to an irrational rotation.

CHAPTER 6 : COUPLED SYSTEMS

OUR INITIAL PATH through dynamics is limited by the restriction to one-dimensional systems. Reality is high-dimensional. As a precursor to more interesting systems, we consider what happens when we couple a pair of 1-D dynamical systems together. For a very restricted class of such coupled systems, it is possible to reduce to a case analyzable via 1-D techniques.

DRIVERS

Consider a pair of agents whose (real) states are increasing at the same speed: imagine two cars driving in separate lanes on the road (where state is position). Model these two states as real variables x_1 and x_2 . If each of these is increasing at the same rate ω , then there is a pair of dynamical systems given by:

$$\begin{array}{ll} Dx_1 = \omega & : \quad Ex_1 = x_1 + \omega \\ Dx_2 = \omega & : \quad Ex_2 = x_2 + \omega \end{array}$$

Continuous time (left) or discrete time (right) systems possess equally plain solutions: $x_i(t) = \omega t + x_i(0)$ or $x_i(n) = \omega n + x_i(0)$ respectively. These two evolving states remain as close as their initial conditions: $x_2 - x_1 = x_2(0) - x_1(0)$.

Now suppose that there is some mechanism of influence: the state of the first influences the second and vice-versa. If, perhaps, the first driver's state x_1 is behind that of the second, x_2 , then the second slows down a bit and the first speeds up. This can be modelled by a function of the *state difference* $\varphi = x_2 - x_1$. Let $f = f(\varphi)$ be a coupling function that delineates the influence of one state on the other. Assuming the effect of the influence is small and symmetric, the appropriate model for the evolution of these two drivers is:

$$\begin{array}{ll} Dx_1 = \omega + \epsilon f(\varphi) & : \quad Ex_1 = x_1 + \omega + \epsilon f(\varphi), \\ Dx_2 = \omega - \epsilon f(\varphi) & : \quad Ex_2 = x_2 + \omega - \epsilon f(\varphi). \end{array}$$

As before, continuous (left) and discrete (right) time systems are similar.

These are *two-dimensional* dynamical systems, not solvable using the tools we have learned thus far. However, because the coupling term depends on the state difference $\varphi = x_2 - x_1$, we can manipulate this system into a reduced system using only the difference variable φ . How does φ evolve? Because the evolution operators D and E are *linear* operators, the derivation is simple:

$$\begin{aligned} D\varphi &= D(x_2 - x_1) = Dx_2 - Dx_1 = (\omega - \epsilon f(\varphi)) - (\omega + \epsilon f(\varphi)) = -2\epsilon f(\varphi), \\ E\varphi &= E(x_2 - x_1) = (x_2 + \omega - \epsilon f(\varphi)) - (x_1 + \omega + \epsilon f(\varphi)) = \varphi - 2\epsilon f(\varphi). \end{aligned}$$

These are in each case a function of φ alone: the 2-D system has been reduced to a 1-D system on the state difference. Now we can use our tools and search for equilibria in φ . In both discrete and continuous time, equilibria occur at roots where $f(\varphi) = 0$: where there is no influence, solutions continue apace at a fixed difference.

SPINNERS

If we consider a pair of agents with *circular* as opposed to real states, the problem of synchronized phase becomes more subtle. Recalling the topology of the circle $\mathbb{S}^1$ from Chapter 5, we choose a pair of circular states, θ_1 and θ_2 , represented as angles modulo 2π . The constant-speed evolution of the form

$$\begin{array}{ll} D\theta_1 = \omega & : \quad E\theta_1 = \theta_1 + \omega \\ D\theta_2 = \omega & : \quad E\theta_2 = \theta_2 + \omega \end{array}$$

exhibit solutions $(\theta(t) = \omega t + \theta_i(0))$ or $\theta_i(n) = \omega n + \theta_i(0)$ respectively) which, instead of moving monotonically, rotate without end. The addition of a small coupling term gives a system of the form:

$$\begin{aligned} D\theta_1 &= \omega + \epsilon f(\varphi) & : & & E\theta_1 &= \theta_1 + \omega + \epsilon f(\varphi) \\ D\theta_2 &= \omega - \epsilon f(\varphi) & : & & E\theta_2 &= \theta_2 + \omega - \epsilon f(\varphi) \end{aligned}$$

Here, $\varphi = \theta_2 - \theta_1$ is the **phase angle**, the angular difference between the two spinners. The coupling function $f: \mathbb{S}^1 \rightarrow \mathbb{R}$ must be chosen with care, as it has the circle as its domain. In particular, if coordinatized by radians, it must be well-defined up to multiples of 2π : periodic functions of φ , such as $\sin \varphi$ or $\cos \varphi$ are permissible, but typical polynomial functions of φ are not.


The derivation of a synchronized state is exactly as it was with drivers, yielding:

$$D\varphi = -2\epsilon f(\varphi) \quad : \quad E\varphi = \varphi - 2\epsilon f(\varphi) \quad (*)$$

With sinusoidal coupling $f(\varphi) = \sin \varphi$, the spinners synchronize, since $\varphi = 0$ is an equilibrium for both systems. The derivative of the right-hand side evaluated at the equilibrium is -2ϵ (continuous) and $1 - 2\epsilon$ (discrete) respectively. This is stable: always in continuous time and when $\epsilon < 1$ in discrete time. *Nota bene*: there is another equilibrium at $\varphi = \pm\pi$, where the spinners are perfectly out-of-phase: this equilibrium is unstable. Such a non-stable equilibrium is nearly unavoidable.


LEMMA: In any spinner system of the form $(*)$ with continuous coupling function $f: \mathbb{S}^1 \rightarrow \mathbb{R}$, there is no stable equilibrium which attracts all initial conditions.

▷ Proof: A stable equilibrium at φ_* implies that $f(\varphi_*) = 0$ and this is a simple zero (f takes on positive and negative values in an arbitrarily small neighborhood of φ_*). Continuity of f on $\mathbb{S}^1$ implies that there is another root of f . This is a distinct equilibrium which cannot be attracted to φ_* . ◁



SYNC vs. LOCK

A stable equilibrium at $\varphi = 0$ is often called synchronization, either in the “drivers” or “spinners” model: both parties converge to the same state. Such convergence can happen in systems with more than two agents, and this can be very satisfying to watch in real-time. It is also the basis of a number of interesting physical, biological, and social phenomena, from flocking to locomotion: see the next subsection.



Synchronization is just the beginning: depending upon the configuration and the coupling function f , one can have stable

equilibria at a nonzero values of φ . This phenomenon is often referred to as (phase) **locking**, especially when the states are angles and the difference φ is a phase angle. For example, if you walk about freely without paying attention, you may notice that your arms are swinging out-of-phase with one another. This is not a coincidence, and you are not purposefully doing it (though you can cease if you think about it). You are flowing along a dynamical system to a stable equilibrium at $\varphi \approx \pm\pi$. In the example of a pair of drivers at positions x_1 , x_2 , it may happen (if they are in the same lane) that there is a stable equilibrium at a nonzero difference $\varphi_* \neq 0$ corresponding to a persistent gap between vehicles (with very unpleasant dynamics if $\varphi_* = 0$). Each of these is an example of a phase lock.

APPLICATIONS

Coupled dynamical systems are central to this text. This initial foray, though narrow in scope, suggests several interesting applications.

PRICE FIXING


When introducing the “drivers” model it was assumed that the agent states x_1 and x_2 corresponded to positions along a line, imagined as vehicles. If instead you view these states as *prices*, then one has a nice model of a pair of competing sellers who set price based on a rate of inflation (assumed constant) with a small degree of correction based on φ , the price difference with the competition. With a locally linear coupling function $f(\varphi)$, then, as per the previous analysis, one tends to a fixed price for both sellers. Note that this suspicious-looking convergence to the exact same price over time takes place without intentional collusion: it is a natural outcome of small adjustments. Such alignments can be seen in data coming from, e.g., private college tuitions in the USA. What could be the cause for persistent price differences in other markets? Is it a lack of mutual influence, or are there other factors?

CENTRAL PATTERN GENERATORS

Biological applications of synchronization are perhaps most satisfying, a field of fireflies in a synchronized light show being a beautiful and well-trod example. Rigorous experiments are easier with a different biological system. A **central pattern generator** (CPG) is a neural network consisting of multiple neuronal “spinners” producing oscillatory outputs without sensory feedback. Such CPGs appear to be fundamental components in locomotion for vertebrates, with conclusive evidence accumulated for lampreys, salamanders, and more. Here, the phenomenon of phase locking dominates. Think of the manner in which an eel swims through a graceful coordination of undulations: these can be modelled by phase-locked CPGs.

BIOMIMETIC ROBOTICS

Coupled oscillators are extremely relevant to locomotion in biological and mechanical agents alike. Ordinary human walking involves two pair of abstract spinners in sync (left-arm-right-leg and right-arm-left-leg), the two sets out-of-phase with each other. Organisms with more limbs (cockroaches, centipedes, millipedes) or more locomotive gaits (trot, gallop, run)



reveal richer sorts of phase locking. In robotics, researchers can program the controls on a multi-legged agent (bipedal, quadrupedal, or hexapedal being common) to engender these or other gaits. This is a rich domain of research, with considerations of symmetry, physics, and phase in play.

SOCIAL NETWORKS

Applications to specific models in biology or robotics are uncontroversial. This changes when the system in question is humans coupled by social interaction. It certainly seems to be the case that a group of individuals reciting a song or a creed do so with a coherence that would be lost without auditory feedback from neighbors. To what degree is it possible to entrain human behaviors or even opinions based on subtle positive or negative reinforcement from peers? Can stability and convergence be analyzed or controlled? This is the subject of interesting though more speculative work in fields ranging from marketing to psychology and more. Specific, well-defined models of opinion dynamics on social networks will be considered in Volume 3.

WHAT IF...

The applications above should prompt a great many questions.

1. What happens if the “natural” rate of change, ω , is not the same for both agents? Do they still synchronize states?
2. The coupling strength ϵ seems not to matter for synchronization. How does changing ϵ impact the system behavior?
3. What happens if there are more than two agents, all influencing each other? Do they still synchronize, and, if so, how quickly?
4. What if there are multiple agents whose network of influence is not all-to-all coupling, but rather local and arranged in a line or a circle?
5. What if the coupling functions differ between agents or change over time?
6. Does directional (one-way) coupling change anything? What if there is an external driver which influences but is uninfluenced by others?

Some of these questions can be answered with the tools we have thus far learned; others are much more difficult and will require additional resources. The first question is one that both is answerable and yet leads to the next principle of this subject.


For concreteness, consider the spinner model in continuous time, where the natural frequencies are given by constants ω_i for $i = 1, 2$. The dynamical system on angles is given by:

$$D\theta_1 = \omega_1 + \epsilon \sin \varphi \quad : \quad D\theta_2 = \omega_2 - \epsilon \sin \varphi$$

The phase difference $\varphi(t) = \theta_1(t) - \theta_2(t)$ has dynamics:

$$D\varphi = D(\theta_2 - \theta_1) = (\omega_2 - \omega_1) - 2\epsilon \sin \varphi$$

This has equilibria where $\sin \varphi = (\omega_2 - \omega_1)/2\epsilon$; this yields a pair of equilibria for sufficiently small values of $\omega_1 - \omega_2$ (or sufficiently large values of ϵ). However, whenever $2\epsilon = |\omega_2 - \omega_1|$, the two equilibria merge into a single degenerate equilibrium (at $\varphi = \pm\pi/2$, depending on sign), and for smaller ϵ , there are no equilibria and no phase-locking phenomena. Only a certain degree of variation in natural frequencies permits phase locking.



This phenomenon of equilibria changing depending on constants is an example of a *bifurcation* – a significant phenomenon to be investigated directly.

CHAPTER 7 : BIFURCATION THEORY

THE NEXT BRANCH in our classification of dynamical phenomena returns to equilibria and the examination of what happens in a degenerate case. The strategy of using linearization to classify equilibria is effective, so long as the linear term in the Taylor expansion does not vanish. One is comforted by that supposition that such is a rare event. In individual dynamical systems, this holds. This chapter introduces parametrized families of dynamical systems, for which failure of linearization is not merely possible but common and classifiable.

PARAMETERIZED SYSTEMS

Let μ denote a *parameter*, by which we mean a variable that a user may change but which time may not. Thus, for autonomous continuous or discrete-time systems, we have, respectively,

$$Dx = f(x, \mu) \quad ; \quad Ex = f(x, \mu) .$$

In both continuous and discrete time settings, μ is a constant ($D\mu = 0$; $E\mu = \mu$ respectively); however, μ can be changed and the dynamical system rerun. The mental image of a multiverse of contingent systems is perhaps helpful.

As per the remarks of Chapter 1, such parametrized dynamical systems are in fact the norm in applied settings, and multiple parameters are common. Knowing which behaviors are possible and how these change as a function of parameter values is crucial for planning and control.

EXAMPLE: A simple continuous-time model for the velocity, v , of a body falling with air resistance is as follows:

$$\frac{dv}{dt} = -g + \kappa v^p ,$$

where $\kappa > 0$ is a parameter known as the *drag coefficient*. The power, p , is an additional parameter, usually chosen to be 1 or 2, depending on the physics of the problem. These are both parameters, since they do not change over time, but can be changed to modify the model.

EXAMPLE: Recall from the previous chapter, the continuous-time model of the phase angle difference between a pair of coupled spinners with different natural frequencies:

$$\frac{d\varphi}{dt} = (\omega_2 - \omega_1) - 2\epsilon \sin \varphi .$$

This model has three parameters: the two frequencies, ω_i , and the coupling strength ϵ . If one fixes values of $\omega_1 \neq \omega_2$ and then runs the system at various coupling strengths, the different behaviors emerge. For $\epsilon \gg 0$ sufficiently large, there are two equilibria; but for sufficiently small $\epsilon > 0$, there are no equilibria. The precise value where things change is at:

$$\epsilon_* = \arcsin \frac{\omega_2 - \omega_1}{2} .$$

Such special parameters are the subject of this chapter.

LOCAL BIFURCATIONS

A (*local*) *bifurcation* of a dynamical system occurs at an equilibrium x_* and parameter μ_* if, in any neighborhood of (x_*, μ_*) , there is a change in the number or types of equilibria. Thanks to the Stability Criterion, bifurcations occur only when the criterion fails.

LEMMA: If a dynamical system of the form $Dx = f(x, \mu)$ or $Ex = f(x, \mu)$ is C^1 (the function f is continuously differentiable in x and μ), then a bifurcation cannot occur unless f vanishes to first order in x .

▷ *Proof:* For the continuous-time case, let f be smooth and $f(a, b) = 0$ with $\partial f / \partial x \neq 0$ at (a, b) . By the Implicit Function Theorem, there exist nearby equilibria of the form $x = x(\mu)$ for (x, μ) near (a, b) . Since f is C^1 , $\partial f / \partial x \neq 0$ along this path of equilibria for nearby μ . This completes the proof for continuous time. In the discrete time case, repeat the argument using $f - x$. ◁

This is more significant than might appear. In a parametrized dynamical system, if you are at an equilibrium and you change the parameter, then, unless there is a bifurcation, the equilibrium may change its location slightly, but its type will remain the same. This holds true in a physical, biological, or social system, even if you do not understand the mechanism by which equilibrium is restored. Thus, if you wish to argue that changing a parameter in a system currently at equilibrium will lead to a cascade of consequences (positive or negative) amounting to a change in stability, then it is tantamount to arguing that there is a bifurcation in the system. The reader is encouraged to look carefully for such arguments in contemporary discussions of economics, climate change, public policy, social dynamics, AI, and more.

NORMAL FORMS OF BIFURCATIONS


Bifurcations arise when the Stability Criterion (*i.e.* linearization) fails. In this case, it is the higher order terms in the Taylor expansion that are determinative. These terms serve as a filtration to order the complexity of bifurcations. In what follows, we restrict to continuous-time settings and provide the simplest possible examples – called *normal forms* – of three elementary bifurcations.

SADDLE NODE BIFURCATIONS [SN]

The first and simplest example of a bifurcation is called a *saddle-node*. It has normal form

$$\frac{dx}{dt} = \mu + cx^2,$$

where $c \neq 0$. The equilibria are located at $x = \pm\sqrt{-\mu/c}$. For $c < 0$, this means



that there are two equilibria when $\mu > 0$; one when $\mu = 0$; and none for $\mu < 0$, with the pattern reversing for $c > 0$. The two equilibria, when they exist, have opposite stability types, as indicated by the derivative $2cx$. The term *saddle-node* is (for now) unmotivated; more sense will be made with the addition of a dimension. For the present, it suffices to remember that saddle-nodes involve a pair of equilibria – one stable, one unstable – colliding and


annihilating.

TRANSCRITICAL BIFURCATIONS [TC]

The next example of a bifurcation is not as common, but it is important. The *transcritical bifurcation* has normal form


$$\frac{dx}{dt} = \mu x + cx^2,$$

where $c \neq 0$. The equilibria are located at $x = 0$ and $x = -\mu/c$. Linearization tells us that the equilibrium at $x = 0$ is stable for $\mu < 0$ and unstable for $\mu > 0$. When $\mu \neq 0$ the other equilibrium at $x = -\mu/c$ has stability opposite that of the equilibrium at zero – unstable for $\mu < 0$ and stable for $\mu > 0$. This bifurcation is characterized by a pair of opposite-stability equilibria which cross each other, exchanging stabilities.



PITCHFORK BIFURCATIONS [PF]


A slight change to the transcritical normal form yields the normal form for a *pitchfork bifurcation*.



$$\frac{dx}{dt} = \mu x + cx^3,$$

where $c \neq 0$. As before, the equilibrium located at $x = 0$ is stable for $\mu < 0$ and unstable for $\mu > 0$. However, there is also a pair of equilibria at $x = \pm\sqrt{-\mu/c}$ when μ and c are of opposite sign. This pair of equilibria have the opposite stability type to that

at $x = 0$: unstable if $\mu < 0$ and stable if $\mu > 0$. The signature of a pitchfork is that a single equilibrium reverses stability as it sheds a symmetric pair of equilibria of the original stability.



SUPERCritical OR SUBCRITICAL?

In all three examples above, the parameter μ is the coefficient of the linear term in x , and c is the coefficient of the dominant higher-order term. In the case of a saddle-node or transcritical bifurcation, the sign of c has marginal impact on the resulting dynamics. However, in the case of the pitchfork bifurcation, the sign of c is critical to the system behavior. When $c < 0$, we say that it is a **supercritical** pitchfork bifurcation [PF*]; for $c > 0$, this is a **subcritical** pitchfork [PF.] (note the super/sub indicated by the accent). The distinction matters a great deal to the overall stability of the system.

Supercritical pitchforks are characterized by a transfer of stability: the stable equilibrium gifts its stability to the pair of new equilibria post-bifurcation. Whatever instability exists is strictly bounded and contained by the pair of stable equilibria. This is reversed for subcritical bifurcations. If the reader prefers, super- or sub-criticality is also indicated by the nonlinear stability of the equilibrium at the bifurcation value.

Why does the difference between super- and sub-criticality matter? Assume that you are (very close to being) at a stable equilibrium and you pass through a pitchfork bifurcation after changing a parameter slightly. You start to notice your near-equilibrium state slipping away, the distance growing exponentially. Perhaps you wait a bit too long before taking action and trying to reverse course. When you turn the dial back to the original parameter, what happens? If your pitchfork was supercritical, the instability is reversible, and you glide back to a stable equilibrium. However, in the subcritical case, you may be too late, and only a drastic over-correction in the parameter holds any hope of

returning you to stability. It is dangerous to overgeneralize details of cartoonish 1-D mathematical models to issues of public policy; yet, the existence (and relative ubiquity) of subcritical pitchforks should give one pause when dismissing slippery slope arguments out of hand.

A FEW MECHANICAL EXAMPLES


It can be difficult to *see* bifurcations or to get an intuition for them. The following simple mechanical examples of the three types of elementary bifurcations in continuous-time may help.

BUCKLING BEAMS : [PF*]

Pitchfork bifurcations are relatively easy to find, so long as there is the correct symmetry. Consider a thin elastic beam with pinned ends under a force of magnitude λ . The beam may deflect with the midpoint being bent away from the axis of the beam by an amount $x(t)$. Assuming that the elastic forces are sufficiently damped, the equation of motion for $x(t)$ is

$$\frac{dx}{dt} = \lambda x - x^3.$$

This has equilibria at $x = 0$ and $x = \pm\sqrt{\lambda}$ for $\lambda > 0$. This means for $\lambda < 0$, the beam is being pulled, and the only equilibrium solution has zero displacement and is stable. However, when $\lambda > 0$, the equilibrium at $x = 0$ becomes unstable, and there are in addition two equilibria at $x = \pm\sqrt{\lambda}$. The preponderance of stable equilibria implies that this is supercritical pitchfork bifurcation at $x = 0, \lambda = 0$.




INVERTED MAGNETIC PENDULUM : [PF.]

Subcritical pitchforks have but a fragile wedge of stable behaviors, and one wishes to not find too many organic examples. One can be constructed by taking an inverted rigid-rod pendulum with a magnetic bob and fixing another magnet above it, with the parameter being the polarity (modeled with intensity as a signed real number). When the magnets are repelling, there is a single unstable equilibrium at the vertical. Changing the magnet to attracting switches this to a stable equilibrium, but with a pair of unstable equilibria to each side, as evidenced by the catastrophic fall

once the magnetic attraction is overcome by the gravitational pull.

SPRING-LOADED INVERTED PENDULUM : [TC]

Transcritical bifurcations require a special type of symmetry that is less common in mechanical systems. Although these appear frequently in population models and models of lasers, neither is especially intuitive. Consider, then, the somewhat convoluted mechanical example illustrated, augmenting the magnetic inverted pendulum with a neutral spring that pushes back from one side. This system experiences a switch in a stable-unstable pair as one changes the suspended magnet from attracting to repelling.




TORQUED PENDULUM : [SN]

Given a rigid-rod pendulum with rather a lot of damping, it may be approximated with a 1-D model on the angle $\theta(t)$ made against the vertical. Such a pendulum has a stable equilibrium at $\theta = 0$ and an unstable equilibrium at $\theta = \pm\pi$. If one applies a constant torque τ to the pendulum, then an appropriate model would be

$$\frac{d\theta}{dt} = \tau - C^2 \sin \theta,$$

where C is a constant and τ is the parameter. This system (do you recognize the equation?) has a pair of equilibria at $\theta = \arcsin \tau/C^2$, so long as $|\tau| < C^2$. For a sufficiently large magnitude of torque there is a saddle-node bifurcation at $\theta = \pm\pi/2$ and $\tau = \pm C^2$. For larger magnitude torques, the pendulum “flips over” and is spun perpetually.




DISCRETE-TIME BIFURCATIONS

The discrete-time versions of the three bifurcations we have seen are a straightforward adaption. A simple substitution of Δ for D yields the following normal forms:

- ▷ SN : $\Delta x = \mu + cx^2 \Rightarrow Ex = (1 + \mu) + cx^2$
- ▷ TC : $\Delta x = \mu x + cx^2 \Rightarrow Ex = (1 + \mu)x + cx^2$
- ▷ PF : $\Delta x = \mu x + cx^3 \Rightarrow Ex = (1 + \mu)x + cx^3$

The interpretations of these bifurcations are the same as their continuous-time counterparts. It is worth looking at the staircase diagrams for each to discern the local changes in number and types of equilibria as μ passes through zero.



Imagining what a system that is close to a bifurcation looks like is an instructive exercise. For example, suppose that a system is very close to a saddle-node bifurcation, but just on the side where there are no equilibria (locally). Despite the absence of equilibria, there is a shadow of the bifurcation to come: iterates slow down and hang about the area before departing. Such approximate equilibria are observable in biological or economic systems in which one remains nearly stationary for some time in an artificial near-equilibrium state.


THE PERIOD-DOUBLING BIFURCATION

There is one new ingredient in discrete-time bifurcations in 1-D. In the above normal forms, the equilibrium at zero changes stability type at $\mu = 0$, where the derivative of the right hand side attains the value $+1$. However, in discrete time, the stability criterion also fails when the derivative evaluates to -1 . What happens when the derivative passes through the value -1 ? Consider the following normal form:

- ▷ PD : $Ex = (-1 + \mu)x + cx^3$

This resembles the pitchfork bifurcation: there is an equilibrium at $x = 0$ which changes its stability type from stable to unstable as the (small) parameter μ goes from positive to negative. Are there any other local equilibria? No, not for small values of x and μ . What happens as μ passes through zero?

A diagram in the case $c > 0$ reveals something new: as μ changes from positive to negative, the stable equilibrium at $x = 0$ becomes unstable, and a small periodic orbit of period 2 emerges. This periodic orbit is itself stable, entraining nearby initial conditions to it. The periodic orbit increases in size with μ : at the bifurcation point, the periodic orbit has collapsed onto the equilibrium.



This case when $c > 0$ is a *supercritical* period-doubling bifurcation. In the case where $c < 0$, the period-2 orbit exists when $\mu > 0$. It is an unstable period-2 orbit, driving nearby initial condition away. As $\mu \rightarrow 0^+$, the unstable period-2 orbit collapses onto the stable equilibrium, rendering it unstable for $\mu < 0$. The reader will note the similarity to the pitchfork bifurcation, though with some confusion of signs. Not for the last time have we seen this type of bifurcation.

CHAPTER 8 : IDENTIFYING BIFURCATIONS

BIFURCATIONS are a beautiful language for classifying how systems change as a function of parameters. For the local bifurcations we have seen so far, there are multiple ways to determine which bifurcations happen where. These explorations will lead us to questions of which bifurcations are most common, and what can occur outside the safety of simple bifurcations.

TAYLOR EXPANSION

Bifurcations arise when the Stability Criterion (*i.e.* linearization) fails. When the linear terms are indeterminate, it is to higher-order terms we must look. Once again, a Taylor series perspective is most useful.

EXAMPLE: Which bifurcation(s) occur at $x = 0$ in the following system?

$$\frac{dx}{dt} = \mu x - (\sin 2x)^2$$

Begin with a Taylor expansion about $x = 0$. This leads to

$$\frac{dx}{dt} = \mu x - \left(2x - \frac{(2x)^3}{3} + O(x^5) \right)^2 = \mu x - 4x^2 + O(x^4)$$

This is a transcritical bifurcation at $\mu = 0$ and $x = 0$. Remember, constants (like the “-4” in from of the x^2 term above) do not change the type of bifurcation.

EXAMPLE: The following experiences bifurcations at $x = 0$:

$$\frac{dx}{dt} = (\mu^2 + 6)x + 5\mu \sinh x$$

What type and at which parameter values is not obvious – a Taylor expansion of the hyperbolic sine is required. To third order, one has:

$$\frac{dx}{dt} = (\mu^2 + 6)x + 5\mu \left(x + \frac{x^3}{6} + O(x^5) \right) = (\mu^2 + \mu + 6)x + \frac{5\mu}{6}x^3 + O(x^5)$$

The linear terms vanish at $\mu = -2, -3$. The form indicates pitchfork bifurcations: which subtype is indicated by the sign of the cubic term at the bifurcation parameter value. No matter whether μ equals -2 or -3 , the cubic term is negative, and both bifurcations are supercritical pitchforks.

EXAMPLE: Consider the following system:

$$\frac{dx}{dt} = \mu x - \frac{x^2}{1+x^2}$$

Begin with a Taylor expansion about $x = 0$. The first term is polynomial; the second is accessible via the geometric series. For x small,

$$\mu x - \frac{x^2}{1+x^2} = \mu x - x^2(1 - x^2 + x^4 + O(x^6)) = \mu x - x^2 + O(x^4)$$

Thus, there is a transcritical bifurcation at $(0,0)$. This is not the only equilibrium. Factoring out an x , we have additional equilibria where

$$\mu = \frac{x}{1+x^2} \Rightarrow \mu x^2 - x + \mu = 0 \Rightarrow x = \frac{1 \pm \sqrt{1-4\mu^2}}{2\mu}$$


This additional pair of equilibria exist only when $|\mu| < 1/2$. When $\mu = \pm 1/2$, there is a single equilibrium at $x = \pm 1$, and this vanishes for $|\mu| > 1/2$. This is a pair of saddle-node bifurcations at $x = \pm 1$ and $\mu = \pm 1/2$.

CONTINUATION

In the previous example, the equilibrium at zero is unstable for $\mu > 0$ by linearization. How would one determine the stability of the two additional equilibria created in the saddle node bifurcation? One is stable and the other unstable, but which is which? Consider the following argument. For $0 < \mu < 1/2$, the two additional equilibria are positive. Since the right-hand-side of the differential equation is a continuous function, the sign of dx/dt alternates between (+) and (-) values except at the saddle-node bifurcation. Thus, the smaller of the two positive equilibria is stable and the larger is unstable: all this is driven by continuity and what is known at $x = 0$.

This type of argument – using continuity to argue an alternation between stable and unstable equilibria – is unique to continuous-time 1-D systems, but it is a very powerful method.

EXAMPLE: The following parametrized system seems difficult to analyze:



$$\frac{dx}{dt} = x(\mu + 1 - x^2)(\mu - 1 + x^2).$$

The equilibria lie in the (x, μ) plane along the line $x = 0$ and the parabola $\mu = \pm(x^2 - 1)$. By continuity, the complement of these curves in the (x, μ) plane consists of connected components on which dx/dt is strictly positive or strictly negative. Knowledge of the sign of dx/dt from a single point in such a region suffices to infer the stability of all the equilibria which border it, since, by continuity, crossing any branch of equilibria (not at a bifurcation) changes

dx/dt from positive to negative or vice versa. For example, when $x \gg \mu$, then dx/dt is negative (since the leading order term is $-x^5$): therefore, the entire boundary of this region in the (x, μ) plane (which lies to the left) has stable equilibria. Crossing any of these curves switches dx/dt from negative to positive: one propagates stabilities accordingly.

In general, you need only one point in the (x, μ) plane in order to determine the sign of $f(x, \mu)$: all stabilities of all branches of equilibria follow from this. In general, good choices for this starting point include places where x or μ are either very small or (better still) very large.

This technique – choosing a far-off point at which evaluation is simple then inferring the rest by continuity – is a type of continuation argument. This method is topological in nature and can be very powerful in more sophisticated settings.


DEGENERACY & CODIMENSION

From the examples we have considered, it would seem that three bifurcations [SN, TC, PF] plus an extra in discrete-time [PD] complete our classification in 1-D. Not so: a great many other bifurcations exist.

EXAMPLE: The following system

$$\frac{dx}{dt} = \mu^2 x - \mu x^2 - \mu x^3 + x^4 = (\mu x - x^2)(\mu - x^2)$$

At first might appear to be a transcritical at $x = 0$, however, when $\mu = 0$, the second and third-order terms also vanish. It is highly degenerate and does not match



any of the bifurcation types yet covered. By factoring the right hand side as $(\mu x - x^2)(\mu - x^2)$ and plotting in the (x, μ) plane, one can view this as a superposition of a TC and SN bifurcation.

Every local bifurcation involves a failure of the Stability Criterion: in continuous-time, this means the first-order term vanishes. What determines the type of bifurcation is the leading-order term(s) in the Taylor expansion. While it is possible that many terms can vanish at the bifurcation point, it would seem to be a “rare” event. What that means is not easy to make precise.

To get at a quantitative notion of degeneracy, consider adding a small perturbation to the right hand side. It is worth contemplating a few examples.

EXAMPLE: Consider what happens when you add a small function of the form $a_0 + a_1x$ (for a_0, a_1 small in absolute value) to the saddle node bifurcation

$$\frac{dx}{dt} = \mu - x^2 + a_0 + a_1x + a_2x = (\mu + a_0) + a_1x - (1 - a_2)x^2$$


No matter what the (sufficiently small) values of a_0, a_1 and a_2 are, this still is a quadratic function and gives a parabola opening downward in the (x, μ) plane. This remains a SN, though at a slightly different location. Note that this would not change even if higher order terms a_kx^k were added to the perturbation. Properly argued, these observations would lead to the conclusion that SN bifurcations are structurally stable as bifurcations.

EXAMPLE: Repeating the above example with a transcritical bifurcation leads to a different outcome. Consider the zero-locus of a TC bifurcation with a 0th order perturbation:

$$\frac{dx}{dt} = a_0 + \mu x - x^2.$$

The behavior of this system depends critically on whether a_0 is positive or negative. For a fixed $a_0 < 0$, the zero-locus is a pair of disjoint parabolae in the (x, μ) plane – one opening up and one opening down, with the two becoming more pointed and morphing into the TC diagram as $a_0 \rightarrow 0^-$. For $a_0 > 0$, the singularity at $(0,0)$ splits the other way, and the zero locus is a pair of disjoint branches which span all (small) values of μ : there are no bifurcations. Again, as $a_0 \rightarrow 0^+$, these two branches touch at the bifurcation point.

One can show (with more work) that adding higher order terms a_kx^k to the above does not change the qualitative features: the TC still splits as above. From this analysis, one concludes that a TC is thought of as a pair of SN bifurcations that meet tangentially. A TC bifurcation is, therefore, slightly more degenerate than a SN: unlike an SN, one can perturb away a TC.



The proper term for describing the degeneracy of a bifurcation is its **codimension**. A detailed definition falls outside the scope of this text, but an intuitive description is feasible, given the examples above. For example:

- ▷ A SN bifurcation has codimension 1: it is persistent in generic 1-parameter families of dynamical systems, but not in 0-parameter families (that is, systems without a parameter).

- ▷ A TC bifurcation has codimension 2: it is persistent in generic 2-parameter families of dynamical systems, but not in 1-parameter families.
- ▷ A PF bifurcation has codimension 2, as will be argued in the next section.
- ▷ A PD bifurcation in discrete time has codimension 1.

Codimension records how many parameters are required for the bifurcation to be persistent (robust with respect to small perturbations). The word *generic* above is doing the heavy lifting.

As per the Lemma at the beginning of Chapter 7, a stable or unstable equilibrium is of codimension 0: it persists under perturbations. From the proof of that Lemma, one guesses that the Implicit Function Theorem (and thus rather a lot of nontrivial analysis) goes into understanding and proving results about codimension.

THE CUSP BIFURCATION

In all the examples of this chapter, there is a single parameter, μ . Such is not the case in more realistic settings, where multiple parameters are common. The following example of a bifurcation in a 2-parameter system helps explain some of the codimension considerations above.

EXAMPLE: The following 2-parameter system,


$$\frac{dx}{dt} = \alpha x \left(1 - \frac{x}{\beta}\right) - \frac{x^2}{1+x^2},$$

is a logistic population model, with a nonlinear correction term. Here, $\alpha, \beta > 0$ are the parameters, controlling growth rates and environmental constraints. This is a model of *infestation*, where population breakouts can occur and stabilize, based on choice of parameter values.

There is clearly an equilibrium at $x = 0$. By linearization, this equilibrium is unstable for all $\alpha > 0$. In the context of a population model, this is natural and uninteresting. What matters are the *stable* equilibria at positive values of x . The nonzero equilibria occur where

$$\alpha \left(1 - \frac{x}{\beta}\right) = \frac{x}{1+x^2} \Rightarrow x^3 - \beta x^2 + \left(1 + \frac{\beta}{\alpha}\right)x - \beta = 0.$$

This is a cubic polynomial in x and gives rise to a 2-D surface in (x, α, β) space as illustrated. For a fixed pair of parameters (α, β) , there are either 1, 2, or 3 real positive solutions to the cubic. By a continuity argument (and the classification at $x = 0$), these equilibria are, in increasing order, stable (if 1), stable – unstable – stable (if 3), or a mix of stable and degenerate (if 2). This serves as a general pattern for an infestation, which has stable population sizes that are small (the infestation is suppressed) or large (the infestation is established).




The graph of such a cubic is called a *cusp singularity* and contains a great deal of structure: an idealized (not to scale) model of a cusp is illustrated. Fixing one parameter (α) and varying the other (β) gives, for most values of α , a branch of stable equilibria along with a stable-unstable branch emanating from a saddle-node. As one tunes α ,

these two sets of equilibria collide in a pitchfork bifurcation. One can thus think of the cusp as an *unfolded pitchfork*.

On the other hand, fixing β and varying α slices the cubic surface in way that, depending on β , either has one branch of stable equilibria, or a “bent” branch with a pair of saddle-nodes as illustrated. The cubic curve of equilibria expresses the phenomenon of *hysteresis*: as one varies α , the stable equilibrium is annihilated in a saddle-node bifurcation, sending one to a far-off stable equilibrium. The salient feature of this hysteresis is that this is irreversible. If one, say, is lax in kitchen cleanliness and a mouse infestation results (one has tuned α past the point of a saddle-node, drifting swiftly to the stable equilibrium of large population size), then a mere return to prior cleanliness (changing α back) does little: one is still stuck at the large-size stable equilibrium. Only a drastic change to force a saddle-node bifurcation can remedy the infestation.

The bifurcation in the previous example is called a *cusp bifurcation*; such is found whenever a pitchfork bifurcation (that normally requires a symmetry to exist) is embedded in a system with an additional (symmetry-breaking) parameter. The cusp bifurcation is an excellent example of a codimension-2 bifurcation. In any cusp bifurcation, one sees in the parameter plane two branches of saddle-node curves emanating from a pitchfork point.

EXAMPLE: One can see (and indeed build) a mechanical example of a cusp bifurcation with some solid pieces and a few springs or rubberbands. The ensuing device (popularized by Zeeman as the “*catastrophe machine*”) has state space $\mathbb{S}^1$ (the rotation angle of a disc) and parameter space $\mathbb{R}^2$ (the location of the end of the second spring in the plane). It is satisfying to feel the hysteretic snap of the saddle-node collapse to the remaining stable equilibrium.



CHAPTER 9 : I-D MYSTERIES

THIS VOLUME on 1-D dynamics has told a simple story. There are two types of systems – continuous and discrete time – with corresponding calculi. In each case, the plan is the same: find the equilibria, linearize the system about the equilibria, then classify their stability. In the case of parameters, one looks for changes in the number or types of equilibria, with Taylor expansion classifying the ensuing bifurcations. The story gets more complicated with the possibility of periodic orbits in discrete-time, or in continuous-time on a circle.

This is the end of the simple story, but it is not the end of dynamics in 1-D. Dynamical systems is a beautiful subject, so much so that indulgence becomes a temptation. The following vignettes are both foreshadowings of more advanced topics in 1-D dynamics and bifurcations, as well as inoculations against excess.

SINGULARITY & CATASTROPHE

There is little room for the numinous in continuous-time 1-D dynamics; nevertheless, when considering bifurcations in systems with multiple parameters, a way emerges. Classification of bifurcations proceeds through understanding Taylor expansions and assigning normal forms and codimension. Both terms have highly technical definitions artfully avoided in this text.

A deeper approach to understanding bifurcations comes from *singularity theory*, which classifies local appearances of the zero-locus of polynomial functions. We have already seen two fundamental types of singularities which (one argues) are persistent under perturbations. These have normal forms as follows:

- ▷ The *fold* singularity: $a + x^2$.
- ▷ The *cusp* singularity: $a + bx + x^3$.

In the first case, one has the 1-parameter unfolding of the SN; in the second, the 2-parameter unfolding of the PF as the point where a pair of SN curves meet in a cusp in the parameter plane. The codimension of these bifurcations matches the dimension of the minimal parameter space needed to unfold it in the normal form.

Why stop there? Singularity theory uses changes of coordinates to produce normal forms of higher-order singularities (and thus higher codimension bifurcations). The next two higher-order singularities have normal forms as follows:

- ▷ The *swallowtail* singularity: $a + bx + cx^2 + x^4$.
- ▷ The *butterfly* singularity: $a + bx + cx^2 + dx^3 + x^5$.

The former has a 3-D parameter space, and the singularity can be visualized; with work, one sees several lower-order bifurcation curves and surfaces in the parameter space. The latter has a 4-D parameter space and is rather more challenging.

The resulting theory and its applications – popularized by Thom in the 1950s and Zeeman in the 1960s – was called *catastrophe theory*. Buoyed by the ubiquity of hysteresis phenomena in all manner of systems – physical, biological, economic, and social – catastrophists conjured swarms of swallowtails in a flurry of articles and books that tipped away from mathematics and more towards socio-morphology. The resulting admixture of mysticism and hype caused the subject to live up to its name: it collapsed.

This is unfortunate, as the mathematics supporting singularity theory branches into elegant and powerful techniques, using *Lie groups* (which obviate the need for ever-more-esoteric names of singularities), *sheaf theory* (and the *intersection homology* which works so well for understanding singularities), and, more generally, *algebraic geometry* (the overarching subject, of which singularity theory is a small corner). Despite the failed hype, Thom and Zeeman were right: bifurcations and singularities surround us.

PERIOD-DOUBLING CASCADES

As noted in Chapter 7, bifurcations in discrete-time 1-D systems are the same as in continuous time, with one exception: the period doubling [PD] bifurcation. This occurs when the value of the derivative at an equilibrium passes through -1 , and these implicate a period-2 orbit, stable or unstable, depending on whether the bifurcation is supercritical [PD*] or subcritical [PD.].

Such is the simple story. What happens when one sees period-doubling in practice can be much more interesting. Consider a logistic model (from Chapter 3) with parameter $\mu > 0$, normalized to represent x as a fraction of population:

$$x_{n+1} = f(x_n, \mu) = \mu x_n(1 - x_n).$$

Any initial condition in $[0,1]$ remains within this interval for all time for any choice of $0 < \mu \leq 4$. This model has equilibria at $x = 0$ and at $x = 1 - 1/\mu$, with the second equilibrium being stable for all $0 < \mu < 3$ (and attracting all initial conditions in $(0,1)$). At $\mu_1 = 3$, there is a supercritical PD* bifurcation, giving rise to a stable period-2 orbit for μ slightly larger than 3. The equilibrium remains unstable for all $\mu > 3$.

What happens next is key. The period-2 orbit can be thought of as an equilibrium of the system with right hand side $f^{(2)} = f \circ f$, as mentioned in Chapter 5. This equilibrium (and thus the period-2 orbit of the original system) can undergo further bifurcations. Indeed, the period-2 orbit remains stable only until $\mu_2 = 1 + \sqrt{6} \approx 3.449$, at which point the period-2 orbit becomes unstable, shedding a stable period-4 orbit. It has undergone another PD*. This period-4 orbit remains attracting for a short window until, at $\mu_3 \approx 3.544$, there is a third PD* to a stable period-8 orbit. Period-16 is born at $\mu_4 \approx 3.564$ and period-32 at $\mu_5 \approx 3.569$. Within a narrow window of parameter space, this cascade of PD's accumulates to a limit

$$\mu_\infty = \lim_{n \rightarrow \infty} \mu_n \approx 3.5696916089 \dots$$

Past this, more chaotic behaviors ensue.

UNIVERSALITY IN PERIOD-DOUBLING

The example of the logistic model undergoing a period-doubling cascade is not exceptional – it is rather the rule. The reader may try using $f = \mu \sin x$ or $f = \mu - x^2$ or any function that has a simple local maximum. In many cases, there appears to be a PD cascade, with the doubling accumulating at some critical, explosive parameter.

Plotting a diagram in the (x, μ) plane as we did with continuous-time systems leads to beautiful and mysterious plots. It is best to plot points by taking some number of initial conditions and plotting the eventual behavior – say, after 1000 or so iterates. One can see both the PD cascade as well as the chaos that

follows. Focusing on the period-doublings, one clearly sees the accumulating sequence.

In the 1970s, it was noted independently by Feigenbaum and by Coullet & Tresser that the rate at which the PDs accumulate in many different 1-D models (and even in certain physical systems) has a limiting ratio,

$$\lambda = \lim_{n \rightarrow \infty} \frac{\mu_{n-1} - \mu_{n-2}}{\mu_n - \mu_{n-1}} \approx 4.669201609103 \dots$$

Although the parameter values at which the PDs occur are entirely model-dependent, their asymptotics appear to be related: it was noticed in numerical experiments that this $\lambda \approx 4.669$ always appears as the exponential rate at which PDs accumulate. Said more precisely, if the n^{th} PD bifurcation occurs at parameter μ_n and these limit to a parameter $\mu_\infty = \lim_n \mu_n$, then the observational evidence is that, for some constant C ,

$$\mu_n - \mu_\infty \sim C\lambda^{-n}.$$

Much ado was made over the discovery of this new “universal” *Feigenbaum number*, with comparisons to π , e , or i , as part of the *Brandnew Science of Chaos*. A mysterious irrational number, together with the fractal-like bifurcation diagram with the potential for infinitely-fine borders between chaos and order is an irresistible temptation for mathematical mysticism.

Such numerological awe is misplaced. This number, λ , is usually given a different symbol in the literature (typically δ or ρ). The argument for using “ λ ” can only be hinted at this early in the text: it is the **dominant eigenvalue** of an operator on 1-d discrete-time dynamical systems. In Volumes 2 and 3, eigenvalues and eigenvectors will become part of our repertoire for understanding dynamical systems. Thanks to the hard work of many mathematicians (especially Lyubich, McMullin, and Sullivan), the experimentally-motivated conjectures about the universality of the asymptotics of period-doublings have precise hypotheses under which precise theorems and proofs hold. That this arises from eigenvalues from a dynamical system on the space of dynamical systems is what is worthy of awe.

CHAOS


The logistic model is indeed chaotic for values of $\mu > 3.57$, as are many other 1-d dynamical systems. What does this mean? The diagrams look “wild”, but something more principled is in order. We give a precise definition of chaos in Volume 4; for now, a foreshadowing is appropriate.

The logistic map at the parameter $\mu = 4$ is, in a sense, more chaotic than at the point $\mu = 3.57$ just past the accumulation of PDs. At $\mu = 4$, orbits of the logistic model appear to bounce about the interval $[0,1]$, independent of the initial condition. This cannot be quite true, as the system clearly has two (unstable) equilibria. Such exceptions pass also to periodic orbits. As we will argue in Volume 4, this system at this parameter value has periodic orbits of **all** periods. They are not readily seen because they are all unstable. They are, not, however, hidden: given any initial condition $x_0 \in [0,1]$ and any $\epsilon > 0$, there are infinitely many periodic points within ϵ of x_0 . This is not obvious.

Is that, then, what chaos entails? Not quite. In 1-D discrete-time systems, one can have chaotic dynamics without having periodic orbits of all periods. In Volume 4, we shall learn of a wonderful theorem that implies, for example, that having a periodic orbit of period 44 forces the existence of periodic orbits

of periods 48 and 32, but not (necessarily) of periods 36, 42, or 11. Chaos in 1-D is well-ordered.

INTERLUDE LINEAR ALGEBRA



CHAPTER 10 : COUPLED & DECOUPLED SYSTEMS

SYSTEMS WITH INTERACTING variables are at the heart of modern applications. Our ability to deal with multiple evolving states, each influencing others, will be entirely dependent upon an update to the mathematical tools used. That the single-variable calculus of Chapter 2 must be upgraded to multi-variable calculus is clear. This ascent is predicated upon a fundamental enhancement in the algebra upon which linear mathematics and its models are built. This concise interlude begins with the dynamical motivations that naturally segue into an introduction to linear algebra.

INTERACTIONS

The next Volume considers dynamical systems with two variables, x and y , both functions of time, continuous or discrete. When the dynamics of one variable depends not only on *its* state but on the *other's* state as well, the system is said to be *coupled*. In Chapter 6, we worked with two types of coupled systems – drivers and spinners – whose tidy analysis comprised a *just-so* story of reducing everything to a 1-D system on the relative differences between states.

This will no longer do. The dynamical systems of the next Volume are of the following general forms:

$$\begin{array}{ll} \frac{dx}{dt} = f(x, y) & : \quad x_{n+1} = f(x_n, y_n) \\ \frac{dy}{dt} = g(x, y) & : \quad y_{n+1} = g(x_n, y_n) \end{array}$$

with left and right specifying continuous and discrete time systems respectively. Such a general case is far too difficult to solve explicitly, and the qualitative methods introduced in Volume 1 must be upgraded to 2-D. Our story remains the same. Equilibria (constant solutions) will be identified, linearized about, and classified based on an investigation of linear systems.

The complexities of 2-D dynamical systems will emerge not from increasingly twisted models or manufactured intricacies, but from simple couplings of simple systems in 1-D.

EXAMPLE: Consider the following uncoupled linear systems in continuous and discrete time:

$$\begin{array}{ll} Dx = 2x & : \quad Ex = 2x \\ Dy = -y/2 & : \quad Ey = -y/2 \end{array}$$

These systems each have a single equilibrium at the origin (0,0). Is this equilibrium stable or unstable? The dynamics along the x -axis are unstable; the dynamics along the y -axis are stable. Such **saddle points** are familiar objects from optimization and game theory; their dynamical emanations will emerge as a frequent feature in 2-D systems.

EXAMPLE: Consider the following twisted variant of the previous example:

$$\begin{array}{ll} Dx = y & : \quad Ex = y \\ Dy = -x & : \quad Ey = -x \end{array}$$

These are coupled, though in a very simple manner. The solutions to these systems are made clear by converting them to a second-order linear equation on a single variable, say, x . Doing so by substitution yields the following:

$$D^2x = -x \quad : \quad E^2x = -x$$

The continuous-time version is a simple harmonic oscillator and has general solution a combination of sines and cosines; the discrete-time version is a time-discretization thereof. Both these systems give rise to oscillatory solutions about the origin. Such

equilibria are called *centers* and will drive much of the narrative to come, from oscillations to periodic orbits and more.

TOWARDS LINEAR ALGEBRA

Our story requires a mathematical preface. The crucial tools come from *linear algebra* (though calculus and Taylor expansion remain relevant). For the reader without much experience in linear algebra, the 2-D case is an excellent starting point. The reader who has been exposed to linear algebra but is unconvinced of its utility may find helpful motivations here.

Generalizing from the previous coupled examples, linear dynamical systems in 2-D (continuous or discrete time) are best cast in the language of matrices:

$$D \begin{pmatrix} x \\ y \end{pmatrix} = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \begin{pmatrix} x \\ y \end{pmatrix} \quad : \quad E \begin{pmatrix} x \\ y \end{pmatrix} = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \begin{pmatrix} x \\ y \end{pmatrix}$$

Such a linear system is uncoupled if and only if the matrix is *diagonal*; that is, all off-diagonal entries vanish. The two examples from the previous subsection – the *saddle* and the *center* – are linear systems whose matrices are, respectively,

$$\text{saddle} : \begin{bmatrix} 2 & 0 \\ 0 & -1/2 \end{bmatrix} \quad : \quad \text{center} : \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$$

The former is uncoupled and implicates a mix of stable and unstable dynamics; the latter, though still coupled, is as simplified as possible and corresponds to a pure oscillatory solution. Both these behaviors are clear from the structure of the matrices, the latter being a rotation by a quarter-turn in the plane.

EXAMPLE: Consider the following linear systems in continuous and discrete time:

$$D \begin{pmatrix} x \\ y \end{pmatrix} = \begin{bmatrix} 3 & -5 \\ 2 & -3 \end{bmatrix} \begin{pmatrix} x \\ y \end{pmatrix} \quad : \quad E \begin{pmatrix} x \\ y \end{pmatrix} = \begin{bmatrix} 3 & -5 \\ 2 & -3 \end{bmatrix} \begin{pmatrix} x \\ y \end{pmatrix}.$$

These are very coupled, and it is not immediately obvious how to proceed. One fruitful approach involves a miraculous change of variables. If we let $u = 2x + y$ and $v = x + y$, then, with sufficient work, one can show that:

$$D \begin{pmatrix} u \\ v \end{pmatrix} = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \begin{pmatrix} u \\ v \end{pmatrix} \quad : \quad E \begin{pmatrix} u \\ v \end{pmatrix} = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \begin{pmatrix} u \\ v \end{pmatrix}.$$

This transformed system is transparent and agrees with the center considered previously. Explicit solutions for u and v can be combined to yield solutions for x and y . The reader who goes to the trouble of doing the *sufficient work* above will discover much of the mechanics of the linear algebra to come.

This strategy of *diagonalization* will motivate the linear algebra we detail in the following chapters. Of note is its apparent limitations: some linear systems (including the center we have already seen) are neither diagonal nor diagonalizable. This reflects the fact that a rotation cannot be assembled from two independent 1-D systems: coupling is critical.

TOWARDS CLASSIFICATION

The particular values that arise on the diagonal of a suitably simplified matrix are of prime importance in linear algebra and dynamics. These *eigenvalues* will fully characterize linear systems.

In 1-D, the Stability Criterion was our principal achievement, distinguishing between stable and unstable equilibria based on linearization. Eigenvalues will be the key to our updating the Stability Criterion and will lead to classification results. The 1-D dichotomy of stable versus unstable equilibria bifurcates and blossoms in 2-D into a taxonomy of equilibrium types: sources, sinks, saddles, spirals, centers and more will enter our dynamical lexicon, all delineated by

eigenvalues and captured through simple matrix properties (notably *trace* and *determinant*).

This eigenvalue approach lines up an attack on nonlinear dynamics in 2-D using the same strategies as in Volume 1. First, one finds the equilibria for a system. By linearizing the dynamical system at an equilibrium and applying knowledge of explicit solutions to linear systems, a local approximation to the full nonlinear system ensues. Eigenvalues provide a quick qualitative check on local phenomena, easily computed and classified.

There are, however, some updates to the story. In 1-D, visualization of dynamics was relatively straightforward. In 2-D, this is still feasible, but only in continuous-time systems: discrete-time systems in 2-D are significantly more difficult to illustrate. Our ability to infer the big picture from local linearizations in 1-D was effortless; that will no longer be the case. Technical questions of when linearization is and is not trustworthy come into play, and we will arrive at a novel result – the *Hartman-Grobman Theorem* – with guarantees on faithful linearizations.

TOWARDS MODELLING

Many of the examples in Volume I – from logistic population models to genetic switches – were necessarily cartoonish and implausible. The situation is a little better with the addition of an extra dimension. An increase in mathematical sophistication coincides with increased fidelity and expressiveness of the models presented. Realistic models will have to wait for the capabilities of Volume 3, but a few plausible situations will be considered here, including the following.

POPULATION MODELS: The growth and decay of a given population is rarely uncoupled from that of other species. The ability to work with coupled systems in 2-D will allow for 2-species models of predatory-prey and competitive types. Population cohort models (dividing a population into two or three fluctuating subgroups) will likewise be analyzable. Though these are still limited in scope, the dynamical phenomena that these simple models reveal will be instructive for later explorations in Volume 3 that work with arbitrary numbers of species or cohorts.

SECOND ORDER MODELS: Linear differential equations and linear recurrence relations of second order (involving quadratic polynomials in the operators D and E for continuous and discrete time respectively) will comprise an initial application of our linear-algebraic techniques. These, then, will permit the analysis of nonlinear second-order models, which have applications ranging from physical to biological and economic systems, the latter arising in market systems with time-delays in price between buyer and seller.

OSCILLATION MODELS: Second order continuous-time systems are particularly useful in working with nonlinear oscillators ranging from mechanical vibrations to nonlinear electrical circuits and more. All of this is driven by the existence of centers, arising from the wonders of imaginary eigenvalues. In the context of parametrized systems, these will lead to bifurcations that, in a very realistic way, model the rise of resonant vibrations or oscillations in systems.

TOWARDS GLOBAL PHENOMENA

Perhaps the biggest challenge in this Volume is the role of global qualitative features in controlling dynamics. Consider the bifurcation theory of Volume 1, in which everything was a local phenomenon contained in a neighborhood of an equilibrium. In this Volume, such *local bifurcations* are also abundant, but

accompanied by a few new entrants. One of these, the *Hopf bifurcation*, implicates periodic orbits, which now can and do exist in abundance.

The detection of periodic orbits in 2-D is subtle, though not completely intractable. Several methods for ascertaining or obstructing the existence of periodic orbits will hold our attention later in Volume 2. The detection and control of non-local phenomena will require non-local mathematics: derivatives and Taylor expansion will give way to integrals and topological methods.

CHAPTER II : LINEAR TRANSFORMATIONS

LINEAR ALGEBRA is both the principal tool and the proper perspective with which to handle interconnected systems. This interlude is a swift self-contained primer for those ideas from linear algebra most useful in solving coupled systems, frequently focusing on the 2-dimensional case relevant to this Volume. It is assumed that the student is familiar (from multivariable calculus) with vectors, matrices, and basic matrix algebra, including computing inverses and determinants of a square matrix.

VECTOR SPACES

Linear algebra is the algebra of vector spaces and linear transformations. A vector space is a collection of objects – abstract vectors – which admit algebraic operations and rules familiar from the use of geometric vectors. The ascent from the Euclidean to the abstract setting is *nontrivial*. The following definition requires careful attention and time to master.

A (real) **vector space** is a collection V of objects (called **vectors**) along with two operations, vector addition (+) and scalar multiplication ($\cdot$) which together satisfy the following properties. For all vectors $\mathbf{u}, \mathbf{v}, \mathbf{w} \in V$ and scalars $a, b \in \mathbb{R}$:

- ▷ Vectors can be added, $\mathbf{u} + \mathbf{v} \in V$ in a manner that is both commutative ($\mathbf{u} + \mathbf{v} = \mathbf{v} + \mathbf{u}$) and associative ($\mathbf{u} + (\mathbf{v} + \mathbf{w}) = (\mathbf{u} + \mathbf{v}) + \mathbf{w}$).
- ▷ There exists a *zero vector*, denoted $\mathbf{0} \in V$, which is an additive identity in that $\mathbf{u} + \mathbf{0} = \mathbf{u}$.
- ▷ Every vector $\mathbf{u} \in V$ has an *inverse*, denoted $-\mathbf{u}$, satisfying $\mathbf{u} + (-\mathbf{u}) = \mathbf{0}$.
- ▷ Scalar multiplication, written $a \cdot \mathbf{u}$ (or $a\mathbf{u}$ for simplicity), gives a vector in V with special cases $0\mathbf{u} = \mathbf{0}$ and $1\mathbf{u} = \mathbf{u}$.
- ▷ Scalar multiplication composes correctly: $a(b\mathbf{u}) = (ab)\mathbf{u}$.
- ▷ Both operations are distributive, meaning that $a(\mathbf{u} + \mathbf{v}) = a\mathbf{u} + a\mathbf{v}$ and $(a + b)\mathbf{u} = a\mathbf{u} + b\mathbf{u}$.

Students of multivariable calculus will be familiar with *Euclidean vector spaces* in which vectors are conflated with points in $\mathbb{R}^n$ having explicit coordinates; and where vector addition is component-wise and scalar multiplication impacts the length of the vectors. These are excellent examples and will animate much of what we do to come. It is worthwhile to check that Euclidean vectors satisfy all the properties above.

The notation used above is meant to reduce surprise. One would guess that the additive inverse $-\mathbf{u}$ is related to the scalar multiplication $(-1)\mathbf{u}$. That this is so follows from a combination of the axioms:

$$\mathbf{u} + (-\mathbf{u}) = \mathbf{0} = 0\mathbf{u} = (1 - 1)\mathbf{u} = 1\mathbf{u} + (-1)\mathbf{u} = \mathbf{u} + (-1)\mathbf{u}.$$

Is this overly formal and redundant? Yes, for Euclidean vectors. However, more abstract vector spaces abound, for which intuition does not suffice.

EXAMPLE: Consider the set of all 2-by-2 matrices with real entries. These can be added and rescaled, operating entry-by-entry. The zero matrix is the additive identity. Note that matrix multiplication does not play a role in the vector space structure – the inverse here is an *additive* inverse: $A \mapsto -A$. The vector space of 2-by-2 matrices is “essentially the same” (or in algebraic language, *isomorphic*) to the Euclidean space $\mathbb{R}^4$, where one presents the vectors in 2-by-2 arrays instead of as stacked column vectors.

EXAMPLE: One simple but powerful class of examples of vector spaces comes from the types of functions one works with in calculus. Consider the set of functions $f: \mathbb{R} \rightarrow \mathbb{R}$. Can one add two such functions together? Yes. Can one rescale a function by a scalar? Yes. Such operations turn this set into a vector space whose additive identity is the zero function and whose inverses take f to $-f$ so that $f + (-f) = 0$.

EXAMPLE: The vector space of functions $f: \mathbb{R} \rightarrow \mathbb{R}$ has a discrete counterpart, mirroring the dichotomy between continuous and discrete calculus or dynamics. Consider the set of sequences $a: \mathbb{N} \rightarrow \mathbb{R}$. Sequences can be added and multiplied by a scalar; the additive inverse and additive identity are the same as in the previous example, sampled at the natural numbers for inputs. This example and the previous are the basis for several methods in signal processing (analog and discrete signals) that use linear algebra to characterize and transform or decompose signals.

EXAMPLE: Consider the set of polynomials in a variable x of degree less than or equal to N . Addition and scalar multiplication preserve this set. The additive identity is the zero function (a polynomial of degree zero) and the additive inverse is multiplication by -1 . Sending $N \rightarrow \infty$ takes this vector space to the vector space of formal power series, familiar from calculus. Compare this with the space of discrete signals from the previous example.

Much more can be done. Spaces of parameters, images, solutions to linear dynamical systems, or subsets of data sets are all excellent contemporary examples of abstract vector spaces.

SUBSPACES & SPAN

The operation of scalar multiplication – rescaling a nonzero vector by a real – has the familiar geometric effect of growing, shrinking, nullifying, or flipping vectors, depending on the scalar value. The algebraic analogue of this geometric action is the span of a vector – the set of all multiples of that vector. In general, for any collection of vectors $S = \{\mathbf{v}_i\}$ in a vector space V , their **span** consists of all possible linear combinations:

$$\text{span}(S) = \sum_i c_i \mathbf{v}_i \quad : \quad c_i \in \mathbb{R}.$$

Such a span always contains the zero-vector. It also is a vector space in and of itself – if two vectors are in $\text{span}(S)$, then so are any linear combination.

In general a **subspace** W of a vector space V is a subset of vectors in V that is closed under vector addition and scalar multiplication: it is itself a vector space. The span of any collection of vectors is a subspace. Any vector space (or subspace) has a **dimension** which can be defined as the minimal number of vectors which span the space.

EXAMPLE: In the standard Euclidean vector space $\mathbb{R}^3$, there are two “trivial” subspaces, the zero-vector $\mathbf{0}$ (a zero-dimensional subspace) and $\mathbb{R}^3$ itself (technically a subspace). The more interesting subspaces are the one-dimensional subspaces (lines through the origin) and the two-dimensional subspaces (planes through the origin).

EXAMPLE: Consider the vector space of real cubic polynomials. The set of all quadratic polynomials (of the form $a_0 + a_1x + a_2x^2$) give a three-dimensional subspace of this four-dimensional vector space, since the sum of quadratic polynomials is still quadratic. In general, the degree- k polynomials form a k -dimensional subspace of the $(n+1)$ -dimensional vector space of all real degree- n polynomials.

A collection of vectors which spans a vector (sub)space and does so minimally (in the sense that they are **linearly independent**) forms a **basis**. The standard basis in Euclidean $\mathbb{R}^n$ is the usual set of unit vectors along the coordinate axes, generalizing the $\hat{i}, \hat{j}, \hat{k}$ vectors from 3-D physics.

Given any basis $\mathcal{B} = \{\mathbf{u}_i\}_{i=1}^n$, a vector $\mathbf{v}$ has a well-defined ordered n -tuple of **coordinates** which determine that vector completely via the linear combination,

$$\mathbf{v} = \begin{pmatrix} c_1 \\ c_2 \\ c_3 \\ \vdots \\ c_n \end{pmatrix} = \sum_{i=1}^n c_i \mathbf{u}_i .$$

When writing out vectors in terms of coordinates, a vertical *column vector* must always be used. In Volume II, all vectors will be in spaces of dimension two; thus, any basis has precisely two (independent) elements. The third Volume will extend to higher dimensions.

EXAMPLE: In the vector space of real cubic polynomials, the vectors $\{1, x, x^2, x^3\}$ form a basis, linear combinations of which form any cubic polynomial with coefficients. This is neither the only possible basis nor the most useful: context is everything.

LINEAR TRANSFORMATIONS

Vector spaces by themselves are impotent. What makes vector spaces useful is the ability to send vectors from one vector space to another respecting the structure of the vector spaces. A **linear transformation** $A: V \rightarrow W$ is an operation that respects vector addition and scalar multiplication operations. The input vector space, V , is called the **domain** and the output vector space W is called the **codomain**.

$$A\left(\sum_i c_i \mathbf{v}_i\right) = \sum_i c_i A(\mathbf{v}_i)$$

When V and W have explicit bases, then A is represented as a matrix with $m = \dim W$ rows and $n = \dim V$ columns. In this case, the transformation A acts on a vector $\mathbf{v} \in V$ via matrix-vector multiplication.

There are certain fundamental subspaces associated with a linear transformation $A: V \rightarrow W$. The **image** of A , $im A$, is the subspace of the codomain given by all vectors arising as outputs of A : all vectors $\mathbf{w} \in W$ of the form $\mathbf{w} = A\mathbf{v}$ for some $\mathbf{v} \in V$. The dimension of this image is called the rank, $rank(A)$. The **kernel**, $ker A$, is the subspace of the domain which is the level-set of zero: all $\mathbf{v} \in V$ such that $A(\mathbf{v}) = \mathbf{0} \in W$. This subspace, sometimes called the **nullspace** of A , has dimension called the **nullity** of A , $null(A)$.

LEMMA: For a linear transformation $A: V \rightarrow W$, the kernel is a subspace of V and the image is a subspace of W .

▷ *Proof:* If $A\mathbf{u} = A\mathbf{v} = \mathbf{0}$, then, by linearity, $A(\mathbf{u} + \mathbf{v}) = \mathbf{0} = A(c\mathbf{v})$ and thus $ker A$ is a subspace. Likewise, $A(\mathbf{u}) + A(\mathbf{v}) = A(\mathbf{u} + \mathbf{v})$, and $cA(\mathbf{u}) = A(c\mathbf{u})$; thus $im(A)$ is a subspace. ◁

Linear transformations can be:

- ▷ INJECTIVE (or **1-to-1**) : distinct inputs are sent to different outputs;
- ▷ SURJECTIVE (or **onto**) : every target output is the image of an input;
- ▷ BIJECTIVE (or an **isomorphism**) : injective and surjective.

Isomorphisms – bijective linear transformations – have matrices which are invertible.

LEMMA: A linear transformation $A: V \rightarrow W$ is injective if and only if $null(A) = 0$.

▷ *Proof:* If $A\mathbf{u} = A\mathbf{v}$, then, by linearity, $A(\mathbf{u} - \mathbf{v}) = \mathbf{0}$ and $\mathbf{u} - \mathbf{v}$ lies in $\ker A$. Thus, if $\text{null}(A) = 0$ is zero, then $\mathbf{u} - \mathbf{v} \in \ker A = \{\mathbf{0}\}$ and $\mathbf{u} = \mathbf{v}$ so that A is injective. If $\text{null}(A) > 0$, then A is not injective, since all vectors in $\ker(A)$ are sent to $\mathbf{0}$. ◁

The following is a fundamental theorem for linear transformations.

THEOREM: For a linear transformation $A: V \rightarrow W$ of finite-dimensional vector spaces $\text{rank}(A) + \text{null}(A) = \dim V$.

TOWARDS FACTORIZATION

Linear algebra concerns itself less with vector spaces and more with linear transformations between vector spaces: it is what you do that matters. At first, a student is focused on matrices, matrix algebra, and various algorithmic aspects to problem-solving such as row reduction or Gram-Schmidt. Only later does the more fundamental algebraic perspective emerge: factorization is the key.

Much in the manner that we wish to decompose dynamical systems into simple uncoupled subsystems, linear algebra operates by factoring operations into core components. Of the many factorizations that exist, the following are especially salient examples. In what follows, A is assumed a square matrix for convenience and consistency with our dynamical applications:

- ▷ **LU DECOMPOSITION** : $A = LU$, where L is lower-triangular and U is upper triangular. This is in essence what one does in Gaussian elimination and row reduction, as it allows for easy back-substitution in a linear system of equations.
- ▷ **QR DECOMPOSITION** : $A = QR$, where Q is orthogonal and R is upper triangular. This is what the Gram-Schmidt algorithm accomplishes.
- ▷ **POLAR DECOMPOSITION** : $A = UP$, where U is orthogonal and P is positive semidefinite.
- ▷ **SINGULAR VALUE DECOMPOSITION** : $A = U\Sigma V^T$, where U and V are both orthogonal and Σ is a diagonal matrix of *singular values*. This splits a linear transformation into a rotation, then a rescaling, followed by another rotation.
- ▷ **JORDAN CANONICAL FORM** : $A = VJV^{-1}$, where V is a change-of-basis and J , the Jordan form of A , is a block diagonal matrix that is “as diagonal as possible”, to be explained in detail in Chapter 25.

There is not enough room in the margins of this text to exposit the various uses of canonical decompositions; these are included here as a guide to the curious, for later attention. The only decomposition that will have significant bearing on dynamical systems is the Jordan form, which seeks to change coordinates in order to make a linear transformation as *decoupled* as possible.

CHAPTER 12 : EIGENVALUES & EIGENVECTORS

LINEAR ALGEBRA is both the principal tool and the proper perspective with which to handle interconnected systems. This chapter is a swift self-contained primer for those ideas from linear algebra most useful in solving coupled systems, frequently focusing on the 2-dimensional case relevant to this Volume.

INVARIANT SUBSPACES

Suppose one has a 2-dimensional dynamical system. In the simplest possible case, the system decomposes. Recall the first example from Chapter 10 – a saddle which evolves according to a diagonal matrix:

$$D \begin{pmatrix} x \\ y \end{pmatrix} = \begin{bmatrix} 2 & 0 \\ 0 & -1/2 \end{bmatrix} \begin{pmatrix} x \\ y \end{pmatrix}$$

This system decouples along the x - and y -axes. Any initial vector evolves in its components independently. Consider, however, what happens in a coupled system for which the coordinates are not (x, y) , but rather some linear combination, such as $u = x + y$ and $v = 2x + 3y$. We can rewrite the system in the new coordinates using basic algebra (or a change-of-basis transformation):

$$\begin{aligned} D \begin{pmatrix} u \\ v \end{pmatrix} &= D \left(\begin{bmatrix} 1 & 1 \\ 2 & 3 \end{bmatrix} \begin{pmatrix} x \\ y \end{pmatrix} \right) = \begin{bmatrix} 1 & 1 \\ 2 & 3 \end{bmatrix} D \begin{pmatrix} x \\ y \end{pmatrix} = \begin{bmatrix} 1 & 1 \\ 2 & 3 \end{bmatrix} \begin{bmatrix} 2 & 0 \\ 0 & -1/2 \end{bmatrix} \begin{pmatrix} x \\ y \end{pmatrix} \\ &= \begin{bmatrix} 1 & 1 \\ 2 & 3 \end{bmatrix} \begin{bmatrix} 2 & 0 \\ 0 & -1/2 \end{bmatrix} \begin{bmatrix} 1 & 1 \\ 2 & 3 \end{bmatrix}^{-1} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{bmatrix} 7 & -5/2 \\ 15 & -11/2 \end{bmatrix} \begin{pmatrix} u \\ v \end{pmatrix} \end{aligned}$$

This system, though it looks ever so complicated compared to the diagonal case, is, in essence, the same dynamics. The old x -axis becomes the 1-D subspace in (u, v) coordinates spanned by the vector $u = 1$ and $v = 2$. Any vector $\mathbf{w}$ in this subspace evolves according to the dynamics $D\mathbf{w} = 2\mathbf{w}$; likewise with the y -axis, now identified with the span of the vector $u = 1$ and $v = 3$.

These two subspaces, respected by their respective dynamics, are examples of what one calls an *invariant subspace*: whatever happens in the subspace stays in the subspace.

What would it mean to have an invariant one-dimensional subspace in a higher-dimensional linear dynamical system of the form $D\mathbf{x} = A\mathbf{x}$ or $E\mathbf{x} = A\mathbf{x}$? Choose a nonzero vector, $\mathbf{v}$, which spans the 1-D subspace; then, the linear dynamical system acts on the span of $\mathbf{v}$ as a 1-D system; thus $A\mathbf{v} = \lambda\mathbf{v}$ for λ some constant. The solution would follow directly from that of Volume 1: for $\mathbf{v}_0$ an initial condition in the subspace, the solution would be $e^{\lambda t}\mathbf{v}_0$ in continuous time and $\lambda^n\mathbf{v}_0$ in discrete time.

This is the dynamical foreshadowing of one of the most important decompositions in linear algebra: the decomposition into *eigenspaces*.

SIMPLE EIGENSPACES

Linear transformations are the core of linear algebra: feedback is the genius of dynamical systems. The linear-algebraic incarnation of feedback is the class of linear transformations from a vector space to itself. These self-targeting linear transformations of the form $A: V \rightarrow V$ bear the intimidating name of *endomorphisms*. In explicit coordinates, they are square matrices. In this text, almost all the linear transformations used in dynamics will be endomorphisms.

We have seen that in dynamical systems, one cares about equilibria: they are states which are invariant under the dynamics. Any endomorphism $A: V \rightarrow V$

also has an equilibrium at the origin: $A(\mathbf{0}) = \mathbf{0}$. However, there are other subspaces of V that are *invariant* under A : any vector in the subspace remains in the subspace under the image of A . The largest such subspace is V itself; the smallest is $\mathbf{0}$. What lies between are *eigenspaces* of A .

A *simple eigenspace* of $A: V \rightarrow V$ is a 1-dimensional subspace E of V that is invariant under A . Otherwise said, for any $\mathbf{v} \in E$, $A\mathbf{v} \in E$. Since E is 1-dimensional, this is equivalent to saying that

$$A\mathbf{v} = \lambda\mathbf{v} \quad : \quad \mathbf{v} \in E$$

for some constant λ . This constant – uniquely defined for E – is called an *eigenvalue*; any $\mathbf{v} \neq \mathbf{0}$ in E is an associated *eigenvector* of A . Eigenvalues tend to be rare (discrete); eigenvectors are never unique, since for any scalar c a rescaled eigenvector $c\mathbf{v}$ satisfies $A(c\mathbf{v}) = c(A\mathbf{v}) = c(\lambda\mathbf{v}) = \lambda(c\mathbf{v})$.

How many eigenvalues does an endomorphism have? In what follows, let us assume that $V = \mathbb{R}^2$ and A is represented as an explicit 2-by-2 matrix. It would *appear*, even in this simple setting, that eigenvalues and eigenspaces may or may not exist. Witness the rotation matrix, also from Chapter 10:

$$J = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}.$$

This linear transformation would seem to have no simple eigenspaces: any 1-D vector subspace is rotated by A by an angle $\pi/2$. We will revisit this example.

Beginning with the fundamental equation $A\mathbf{v} = \lambda\mathbf{v}$, move everything to one side and factor out the linear operator like so:

$$(A - \lambda I)\mathbf{v} = \mathbf{0}.$$

Eigenvalues are precisely those values of λ at which the above is satisfied for a nonzero (eigen)vector $\mathbf{v}$. Since an endomorphism is noninvertible if and only if it has nonzero kernel, we conclude that λ is an eigenvalue of A if and only if

$$\det(A - \lambda I) = 0.$$

This is a polynomial equation in λ of degree n , where A is n -by- n : it is the *characteristic polynomial* of A . From this more algebraic perspective, one concludes that a 2-by-2 matrix has a *pair* of eigenvalues, characterized by one of the following three cases:

- ▷ There are two real, distinct eigenvalues $\lambda_1 \neq \lambda_2$;
- ▷ There are two real, repeated eigenvalues $\lambda_{1,2} = \lambda$; or
- ▷ There is a conjugate pair of eigenvalues $\lambda_{1,2} = \alpha \pm i\beta$, with $\beta \neq 0$.

These three cases present a trichotomy that will recur throughout this text.

EXAMPLE: Diagonal matrices are simple. The diagonal entries (even if zero) are precisely the eigenvalues. The x -axis is the eigenspace of the first eigenvalue; the y -axis is the eigenspace of the second, even if the two eigenvalues are the same.

EXAMPLE: Consider the following linear transformation of $\mathbb{R}^2$:

$$A = \begin{bmatrix} 2 & 2 \\ 1 & 3 \end{bmatrix}.$$

The characteristic polynomial is $(2 - \lambda)(3 - \lambda) - 2 = \lambda^2 - 5\lambda + 4$, which has roots $\lambda_1 = 1$ and $\lambda_2 = 4$. Eigenvectors can be chosen by solving $(A - \lambda I)\mathbf{v} = \mathbf{0}$ for $\mathbf{v} \neq \mathbf{0}$. Thus,

$$(A - \lambda_1)\mathbf{v}_1 = \mathbf{0} \Rightarrow \begin{bmatrix} 2-1 & 2 \\ 1 & 3-1 \end{bmatrix} \mathbf{v}_1 = \mathbf{0} \Rightarrow \mathbf{v}_1 = \begin{pmatrix} 2 \\ -1 \end{pmatrix},$$

$$(A - \lambda_2)v_2 = \mathbf{0} \Rightarrow \begin{bmatrix} 2-4 & 2 \\ 1 & 3-4 \end{bmatrix} v_2 = \mathbf{0} \Rightarrow v_2 = \begin{pmatrix} 1 \\ 1 \end{pmatrix}.$$

The case of real, distinct eigenvalues is most desirable, as the eigenvectors form a basis, an **eigenbasis**:

LEMMA: If a linear transformation $A: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ has real distinct eigenvalues, then the associated eigenvectors form a basis for $\mathbb{R}^2$.

▷ *Proof*: By way of contradiction, assume the two eigenvectors are not linearly independent. Then they are parallel and must therefore have the same eigenvalue, belonging to the same eigenspace. ◁

REPEATED EIGENVALUES

In the case of a repeated real eigenvalue λ , there are two possibilities, distinguished by the dimension of the eigenspace. If $\ker(A - \lambda I)$ has dimension two, then, since the vector space is 2-dimensional, every nonzero vector is an eigenvector. This is not the only possibility.

EXAMPLE: Consider the following linear transformation of $\mathbb{R}^2$:

$$A = \begin{bmatrix} 1 & -1 \\ 1 & 3 \end{bmatrix}.$$

The characteristic polynomial is $(1 - \lambda)(3 - \lambda) + 1 = \lambda^2 - 4\lambda + 4$, which has roots $\lambda_{1,2} = 2$. Solving $(A - \lambda I)v = \mathbf{0}$ reveals a 1-dimensional eigenspace, since

$$(A - 2I) = \begin{bmatrix} -1 & -1 \\ 1 & 1 \end{bmatrix}$$

has 1-dimensional kernel spanned by the obvious choice of eigenvector $v = (1, -1)^T$.

A basis of eigenvectors is most desirable for solving linear dynamical systems. In the case of a repeated eigenvalue with 1-dimensional eigenspace, there is a “best” choice for a basis adapted to the linear transformation. Given A with repeated eigenvalue λ and 1-D eigenspace spanned by v , let w be a vector satisfying

$$(A - \lambda I)w = v.$$

Such a **generalized eigenvector**, w , is clearly not parallel to v and thus contributes to a basis. This definition will be of use to us in subsequent chapters.

COMPLEX EIGENVALUES

The final possibility is that one has non-real eigenvalues. This is both simpler and more complex than the repeated-root case just considered. Simpler, because there is no need for generalized eigenvectors; more complex, in that the eigenvectors, like the eigenvalues, come in a complex conjugate pair.

EXAMPLE: Recall the simple rotation matrix J ,

$$J = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}.$$

The characteristic polynomial is $\lambda^2 + 1$, which has roots $\lambda = \pm i$. The eigenvectors can be found as follows:

$$(J - iI)v = \mathbf{0} \Rightarrow \begin{bmatrix} -i & -1 \\ 1 & -i \end{bmatrix} v = \mathbf{0} \Rightarrow v = \begin{pmatrix} i \\ 1 \end{pmatrix}.$$


The second eigenvector is the complex conjugate of the first, reflecting the second eigenvalue.

Students often dislike complex eigenvalues and eigenvectors. They are unpleasant to compute, and the etymological baggage – *imaginary* and *complex* –

evoke a sense of futile burden. The remedy is to be found in future chapters, wherein the intuition of imaginary eigenvalues as characterizing rotations and complex eigenvalues arising from vibrations is developed. The foreshadowing of such is indicated in the matrix J above.

VOLUME 2

THE NEXT DIMENSION



CHAPTER 13 : 2-D LINEAR DYNAMICS

THE ASCENT from 1-D to 2-D dynamical systems impacts everything thus far learned – from classification of equilibria to the Stability Criterion, bifurcations, periodic orbits, and more. The plane is the perfect domain for learning dynamics: it is just complex enough to fascinate, but not so byzantine as to bewilder. With our introduction to linear algebra in-hand, we can begin the story of 2-D dynamics with the linear case.

MATRIX SOLUTIONS

In principle, linear dynamical systems in discrete or continuous time are as easy to solve in arbitrary dimensions as they are in 2-D. We therefore begin with a general solution.

DISCRETE TIME: a discrete-time linear system on $\mathbb{R}^k$ is a dynamical system of the form $E\mathbf{x} = A\mathbf{x}$, where $\mathbf{x} \in \mathbb{R}^k$ is a vector variable. This has an obvious explicit solution in terms of the matrix A and an initial condition $\mathbf{x}_0 = \mathbf{x}(0)$:

$$\mathbf{x}(n) = A^n \mathbf{x}_0 .$$

This is deceptively simple. It motivates the problem of computing arbitrary powers of a square matrix.

CONTINUOUS TIME: a continuous-time linear system on $\mathbb{R}^k$ is a dynamical system of the form $D\mathbf{x} = A\mathbf{x}$ with initial condition $\mathbf{x}_0 = \mathbf{x}(0)$. This first-order ordinary differential equation on $\mathbf{x} \in \mathbb{R}^k$ has solution:

$$\mathbf{x}(t) = e^{At} \mathbf{x}_0 ,$$

where e^{At} is the *matrix exponential*,

$$e^{At} = \exp(At) = \sum_{n=0}^{\infty} \frac{(At)^n}{n!} = I + At + \frac{1}{2!} A^2 t^2 + \frac{1}{3!} A^3 t^3 + \dots$$

LEMMA: $\mathbf{x}(t) = e^{At}$ solves the linear system $D\mathbf{x} = A\mathbf{x}$.

▷ *Proof:* The proof is by direct computation of $D(e^{At})$:

$$\begin{aligned} &= \frac{d}{dt} \sum_{n=0}^{\infty} \frac{(At)^n}{n!} = \sum_{n=0}^{\infty} \frac{d}{dt} \frac{(At)^n}{n!} = \sum_{n=1}^{\infty} \frac{nA(At)^{n-1}}{n!} = \sum_{n=1}^{\infty} \frac{A(At)^{n-1}}{(n-1)!} = A \sum_{m=0}^{\infty} \frac{(At)^m}{m!} \\ &= Ae^{At}. \end{aligned}$$

Note the care in indexing and the use of absolute convergence to differentiate term by term. ◁

This lemma is satisfying to a mathematician. To an engineer or scientist who wants to compute an explicit answer, the problem is but half solved.

EXAMPLE: Diagonal matrices are trivial to take powers of or to exponentiate:

$$\begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix}^n = \begin{bmatrix} (\lambda_1)^n & 0 \\ 0 & (\lambda_2)^n \end{bmatrix} \quad : \quad \exp\left(\begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} t\right) = \begin{bmatrix} e^{\lambda_1 t} & 0 \\ 0 & e^{\lambda_2 t} \end{bmatrix}.$$

The former is trivial and the latter follows from the former and the definition of e^{At} .

This example is the key to computing general solutions to linear systems. If we can diagonalize the matrix – *if we can decouple the system* – then the 1-D solutions suffice.

SIMPLE EIGENVALUES

In 2-D, the solutions to linear systems speciate according to whether the eigenvalues of the associated matrix are real or complex, distinct or repeated. The

simplest case – real, distinct eigenvalues – is only slightly more complicated than the decoupled diagonal case in the previous example.

Given a 2-by-2 matrix A with real eigenvalues λ_1, λ_2 and corresponding eigenvectors $\mathbf{v}_1, \mathbf{v}_2$, the pair of eigenvectors are linearly independent and thus form a basis – an **eigenbasis**. Changing coordinates via the eigenbasis decouples this system. In linear algebra, one learns this as *diagonalization*. In this setting one has:

$$AV = V\Lambda \quad : \quad V = [\mathbf{v}_1 \mid \mathbf{v}_2] \quad : \quad \Lambda = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix}.$$

Note that this follows directly from translating the definition of eigenvalues and eigenvectors into the language of matrices. Powers of A are computed as

$$A^n = (V\Lambda V^{-1})^n = (V\Lambda V^{-1})(V\Lambda V^{-1}) \dots (V\Lambda V^{-1}) = V\Lambda^n V^{-1}.$$

Given that Λ is diagonal, A^n is explicitly computed. This *conjugation* of A^n by the matrix V is precisely a coordinate change via the eigenbasis $\{\mathbf{v}_1, \mathbf{v}_2\}$.

Exponentiating A is only slightly more involved. By the definition of e^{At} , we have

$$e^{At} = \sum_{n=0}^{\infty} \frac{(At)^n}{n!} = \sum_{n=0}^{\infty} \frac{V\Lambda^n V^{-1} t^n}{n!} = V \left(\sum_{n=0}^{\infty} \frac{(\Lambda t)^n}{n!} \right) V^{-1} = V e^{\Lambda t} V^{-1},$$

with some care and justification in the order of matrix multiplication. Knowing what the exponential of a diagonal matrix is, the problem is solved.

EXAMPLE: Given the matrix with eigendecomposition

$$A = \begin{bmatrix} 2 & 5 \\ 3 & 4 \end{bmatrix} \quad : \quad \lambda_1 = 7; \mathbf{v}_1 = \begin{pmatrix} 1 \\ 1 \end{pmatrix} \quad : \quad \lambda_2 = -1; \mathbf{v}_2 = \begin{pmatrix} 5 \\ -3 \end{pmatrix},$$

the solutions to $E\mathbf{x} = A\mathbf{x}$ and $D\mathbf{x} = A\mathbf{x}$ are, respectively,

$$\begin{aligned} A^n \mathbf{x}_0 &= \begin{bmatrix} 1 & 5 \\ 1 & -3 \end{bmatrix} \begin{bmatrix} 7^n & 0 \\ 0 & (-1)^n \end{bmatrix} \begin{bmatrix} 1 & 5 \\ 1 & -3 \end{bmatrix}^{-1} \begin{pmatrix} x_0 \\ y_0 \end{pmatrix} = \frac{1}{8} \begin{pmatrix} (3x_0 + 5y_0)7^n + 5(x_0 - y_0)(-1)^n \\ (3x_0 + 5y_0)7^n - 3(x_0 - y_0)(-1)^n \end{pmatrix}, \\ e^{At} \mathbf{x}_0 &= \begin{bmatrix} 1 & 5 \\ 1 & -3 \end{bmatrix} \begin{bmatrix} e^{7t} & 0 \\ 0 & e^{-t} \end{bmatrix} \begin{bmatrix} 1 & 5 \\ 1 & -3 \end{bmatrix}^{-1} \begin{pmatrix} x_0 \\ y_0 \end{pmatrix} = \frac{1}{8} \begin{pmatrix} (3x_0 + 5y_0)e^{7t} + 5(x_0 - y_0)e^{-t} \\ (3x_0 + 5y_0)e^{7t} - 3(x_0 - y_0)e^{-t} \end{pmatrix}. \end{aligned}$$

DOMINANCE

An examination of the previous example points to a phenomenon that we will repeatedly exploit in understanding linear systems. In that example the eigenvalues are $\lambda_1 = 7$ and $\lambda_2 = -1$: in the continuous-time setting, this means one unstable and one stable eigenvalue. Over time, the solutions to $D\mathbf{x} = A\mathbf{x}$ have a component that rapidly grows (along the unstable eigenspace) and a part that rapidly shrinks to zero (along the stable eigenspace). In the limit as $t \rightarrow +\infty$, the solution asymptotes to

$$\mathbf{x}(t) \rightarrow C e^{7t} \begin{pmatrix} 1 \\ 1 \end{pmatrix} \quad : \quad t \rightarrow +\infty,$$

where the constant C is dependent on the initial condition. Asymptotically, only the first eigenvalue and eigenvector contribute to the solution characteristics.

This is an example of a **dominant eigenvalue**. For the time being, we say that in a continuous-time 2-D system, λ_1 **dominates** λ_2 if $\lambda_1 > \lambda_2$. In discrete time, the condition becomes $|\lambda_1| > |\lambda_2|$, in keeping with the Stability Criterion.

If λ_1 dominates λ_2 , then solutions to a linear system have the following asymptotics:

$$\mathbf{x}(n) = C_1 \lambda_1^n \mathbf{v}_1 + C_2 \lambda_2^n \mathbf{v}_2 = \lambda_1^n \left(C_1 \mathbf{v}_1 + C_2 \left(\frac{\lambda_2}{\lambda_1} \right)^n \mathbf{v}_2 \right) \longrightarrow C_1 \lambda_1^n \mathbf{v}_1$$

$$\mathbf{x}(t) = C_1 e^{\lambda_1 t} \mathbf{v}_1 + C_2 e^{\lambda_2 t} \mathbf{v}_2 = e^{\lambda_1 t} (C_1 \mathbf{v}_1 + C_2 e^{(\lambda_2 - \lambda_1)t} \mathbf{v}_2) \longrightarrow C_1 e^{\lambda_1 t} \mathbf{v}_1$$

The dominant eigenvalue-eigenvector pair controls the long-term behavior of almost all solutions (note the presence of the constant C_1 , which can vanish for certain initial conditions!). In general,

- ▷ The asymptotic growth rate is determined by the dominant eigenvalue.
- ▷ The asymptotic ratio of x to y is determined by the dominant eigenvector.

The following example demonstrates that even when both eigenvalues are unstable (solutions are growing over time), only the dominant eigenvalue matters long-term.

EXAMPLE: Assume two competing Companies with number of customers, A_n and B_n , as a function of year n , is modeled as a linear system. Each year A-customers naturally increase by 7%; for B-customers, this growth rate is 4%. However, each year, 2% of A-customers switch to B, and 1% of B-customers switch to A. The resulting model is:

$$\begin{pmatrix} A_{n+1} \\ B_{n+1} \end{pmatrix} = \begin{pmatrix} A_n + 0.07A_n - 0.02A_n + 0.01B_n \\ B_n + 0.04B_n - 0.01B_n + 0.02A_n \end{pmatrix} = \begin{bmatrix} 1.05 & 0.01 \\ 0.02 & 1.03 \end{bmatrix} \begin{pmatrix} A_n \\ B_n \end{pmatrix}$$

The eigenvalues and eigenvectors of this matrix are:

$$A = \begin{bmatrix} 1.05 & -0.02 \\ 0.01 & 1.03 \end{bmatrix} \Rightarrow \lambda_1 \approx 1.0573; \mathbf{v}_1 \approx \begin{pmatrix} 1.366 \\ 1 \end{pmatrix} : \lambda_2 \approx 1.0227; \mathbf{v}_2 \approx \begin{pmatrix} -0.366 \\ 1 \end{pmatrix}.$$

The long-term dynamics of this system are given by

$$\begin{pmatrix} A_n \\ B_n \end{pmatrix} = \lambda_1^n \left(C_1 \mathbf{v}_1 + C_2 \left(\frac{\lambda_2}{\lambda_1} \right)^n \mathbf{v}_2 \right) \longrightarrow \approx C_1 (1.0573)^n \begin{pmatrix} 1.366 \\ 1 \end{pmatrix},$$

so that, over time, the annual growth rate of both companies is identical at $\approx 5.73\%$. The customer base evolves so that Company A has $\approx 57.7\%$ market share and Company B the remaining 42.3%.

REPEATED EIGENVALUES

The case of repeated real eigenvalues is subtle, as there may be one or two linearly independent eigenvectors. Assume that A is a 2-by-2 matrix with double eigenvalue λ . If the dimension of the kernel of $A - \lambda I$ is two, then an eigenbasis exists, and the results of the previous section suffice to compute powers and exponentials of A : the matrix A is diagonalizable.

In the case that the eigenspace of λ is 1-dimensional, then A cannot be diagonalized, but it can be simplified. Given an eigenvector $\mathbf{v}$, recall from the previous section the generalized eigenvector paired to $\mathbf{v}$, $\mathbf{w}$, satisfying $(A - \lambda I)\mathbf{w} = \mathbf{v}$. Although one cannot diagonalize the matrix A , one can come close.

LEMMA: With $A, \lambda, \mathbf{v}$, and $\mathbf{w}$ as above,

$$A \begin{bmatrix} \mathbf{v} & \mathbf{w} \end{bmatrix} = \begin{bmatrix} \mathbf{v} & \mathbf{w} \end{bmatrix} \begin{bmatrix} \lambda & 1 \\ 0 & \lambda \end{bmatrix}.$$

▷ *Proof:* By direct computation, since $A\mathbf{v} = \lambda\mathbf{v}$, and $A\mathbf{w} = \mathbf{v} + \lambda\mathbf{w}$. ◁

By coming as close to a diagonal matrix as possible, we can solve linear systems and then change coordinates. In what follows, assume $\lambda, \mathbf{v}$, and $\mathbf{w}$ as in the previous Lemma. Consider the following odd-looking matrix:

$$N = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}.$$

This is an example of a **nilpotent** matrix, since some finite power of it is zero: in fact, $N^2 = 0$. That assists in solving repeated eigenvalue linear systems like so:

LEMMA: For $B = \lambda I + N$ as above,

$$B^n = \begin{bmatrix} \lambda & 1 \\ 0 & \lambda \end{bmatrix}^n = \begin{bmatrix} \lambda^n & n\lambda^{n-1} \\ 0 & \lambda^n \end{bmatrix} \quad : \quad e^{Bt} = \exp\left(\begin{bmatrix} \lambda & 1 \\ 0 & \lambda \end{bmatrix}t\right) = \begin{bmatrix} e^{\lambda t} & te^{\lambda t} \\ 0 & e^{\lambda t} \end{bmatrix}.$$

▷ *Proof:* For powers of B , the Binomial Theorem implies

$$(\lambda I + N)^n = \sum_{k=0}^n \binom{n}{k} (\lambda I)^{n-k} N^k = \lambda^n I + n\lambda^{n-1} N.$$

To compute the exponential, note that $e^{(\lambda I + N)t}$ splits as $e^{\lambda t I} e^{Nt}$, though one must be careful with commutativity (the identity matrix helps here...). The definition of matrix exponentiation together with the nilpotence of N implies

$$e^{Nt} = I + Nt = \begin{bmatrix} 1 & t \\ 0 & 1 \end{bmatrix} \Rightarrow e^{At} = e^{\lambda t I} e^{Nt} = \begin{bmatrix} e^{\lambda t} & te^{\lambda t} \\ 0 & e^{\lambda t} \end{bmatrix}. \quad \triangleleft$$

One uses the change-of-basis matrix $V = [\mathbf{v} \quad \mathbf{w}]$ to transform to (generalized) eigenbasis coordinates, importing the solutions to the linear systems above.

EXAMPLE: Recalling from the last chapter the repeated-eigenvalue example,

$$A = \begin{bmatrix} 1 & -1 \\ 1 & 3 \end{bmatrix}.$$

The eigenvalues are $\lambda_{1,2} = 2$ with 1-dimensional eigenspace spanned by $\mathbf{v} = (1, -1)^T$. The generalized eigenvector paired to $\mathbf{v}$ is $\mathbf{w} = (-1, 0)^T$. To solve the linear systems $E\mathbf{x} = A\mathbf{x}$ and $D\mathbf{x} = A\mathbf{x}$ we compute

$$\begin{aligned} \begin{bmatrix} 1 & -1 \\ 1 & 3 \end{bmatrix}^n &= \begin{bmatrix} 1 & 0 \\ -1 & -1 \end{bmatrix} \begin{bmatrix} 2 & 1 \\ 0 & 2 \end{bmatrix}^n \begin{bmatrix} 1 & 0 \\ -1 & -1 \end{bmatrix}^{-1} = \begin{bmatrix} 1 & 0 \\ -1 & -1 \end{bmatrix} \begin{bmatrix} 2^n & n2^{n-1} \\ 0 & 2^n \end{bmatrix} \begin{bmatrix} 1 & 0 \\ -1 & -1 \end{bmatrix}^{-1}, \\ \exp\left(\begin{bmatrix} 1 & -1 \\ 1 & 3 \end{bmatrix}t\right) &= \begin{bmatrix} 1 & 0 \\ -1 & -1 \end{bmatrix} \exp\left(\begin{bmatrix} 2 & 1 \\ 0 & 2 \end{bmatrix}t\right) \begin{bmatrix} 1 & 0 \\ -1 & -1 \end{bmatrix}^{-1} \\ &= \begin{bmatrix} 1 & 0 \\ -1 & -1 \end{bmatrix} \begin{bmatrix} e^{2t} & te^{2t} \\ 0 & e^{2t} \end{bmatrix} \begin{bmatrix} 1 & 0 \\ -1 & -1 \end{bmatrix}^{-1}. \end{aligned}$$

Writing out these matrix multiplications is unpleasant but sufficient.

COMPLEX EIGENVALUES

We have seen that not all matrices can be diagonalized. In the case of complex conjugate eigenvalues, there is a conflict in perspectives. For the mathematician, diagonalization in this case is trivial – over the field of complex numbers. For A with eigenvalues $\lambda, \bar{\lambda} = \alpha \pm i\beta$ and *complex conjugate* eigenvectors $\mathbf{v}, \bar{\mathbf{v}}$, one can assert:

$$V^{-1}AV = [\mathbf{v} \quad \bar{\mathbf{v}}]^{-1} A [\mathbf{v} \quad \bar{\mathbf{v}}] = \begin{bmatrix} \lambda & 0 \\ 0 & \bar{\lambda} \end{bmatrix}.$$

This is unsatisfying to the scientists or engineer, as the solutions to a real linear system should be real-valued functions of time. To adapt, one can take the real and imaginary components of the above solution. Given a complex eigenvector

$$\mathbf{v} = \begin{pmatrix} v_1 + i w_1 \\ v_2 + i w_2 \end{pmatrix},$$

the real and imaginary components together form a change of basis matrix, V ,

$$V = \begin{bmatrix} v_1 & w_1 \\ v_2 & w_2 \end{bmatrix} \Rightarrow V^{-1}AV = \begin{bmatrix} v_1 & w_1 \\ v_2 & w_2 \end{bmatrix}^{-1} A \begin{bmatrix} v_1 & w_1 \\ v_2 & w_2 \end{bmatrix} = \begin{bmatrix} \alpha & -\beta \\ \beta & \alpha \end{bmatrix} = \alpha I + \beta J.$$

Having reduced to the simplest case (the *normal form*), we will find it useful to invoke the matrix version of *Euler's Formula*.

LEMMA: $e^{Jt} = (\cos t)I + (\sin t)J$.

▷ *Proof*: The proof is by direct computation, combined with the fact that powers of J are either $\pm I$ or $\pm J$ based on the power mod 4. Splitting the exponential up into even and odd powers makes the result transparent:

$$\begin{aligned} e^{Jt} &= \sum_{n=0}^{\infty} \frac{(Jt)^n}{n!} = \sum_{n=0}^{\infty} \frac{(Jt)^{2n}}{(2n)!} + \sum_{n=0}^{\infty} \frac{(Jt)^{2n+1}}{(2n+1)!} \\ &= \sum_{n=0}^{\infty} (-1)^n I \frac{t^{2n}}{(2n)!} + \sum_{n=0}^{\infty} (-1)^n J \frac{t^{2n+1}}{(2n+1)!}. \end{aligned}$$

Recalling the Taylor series for sine and cosine completes the proof. ◁

COROLLARY: For eigenvalues $\lambda_{1,2} = \alpha \pm i\beta$, the exponential of $(\alpha I + \beta J)t$ is:

$$e^{(\alpha I + \beta J)t} = \exp\left(\begin{bmatrix} \alpha & -\beta \\ \beta & \alpha \end{bmatrix} t\right) = \begin{bmatrix} e^{\alpha t} \cos \beta t & -e^{\alpha t} \sin \beta t \\ e^{\alpha t} \sin \beta t & e^{\alpha t} \cos \beta t \end{bmatrix}.$$

▷ *Proof*: By commutativity of multiplication with I and the previous lemma,

$$\exp((\alpha I + \beta J)t) = \exp(\alpha t I) \cdot \exp(\beta t J) = (e^{\alpha t} \cos \beta t)I + (e^{\alpha t} \sin \beta t)J. \quad \triangleleft$$

To take powers of the matrix $(\alpha I + \beta J)$, the polar representation helps. Rewriting $\alpha \pm i\beta = re^{\pm i\theta}$, where the modulus, r , and argument, θ , are defined by the usual polar coordinate formulae as

$$r = \sqrt{\alpha^2 + \beta^2} \quad ; \quad \theta = \arctan \frac{\beta}{\alpha},$$

we compute powers of $(\alpha I + \beta J) = (rI)e^{J\theta}$ as

$$\begin{bmatrix} \alpha & -\beta \\ \beta & \alpha \end{bmatrix}^n = (rI e^{J\theta})^n = r^n \begin{bmatrix} \cos n\theta & -\sin n\theta \\ \sin n\theta & \cos n\theta \end{bmatrix}.$$

This is the simplest possible form of a real matrix with complex eigenvalues. The general case, including a change of coordinates to $AV = V(\alpha I + \beta J)$, is computable, but generally unpleasant.

CHAPTER 14 : 2ND ORDER LINEAR SYSTEMS

MATRICES & MATRIX ALGEBRA, while vital, are not the complete picture. Much of the motivation for 2-D dynamical systems comes from linear differential equations or recurrence relations of 2nd order. In this chapter, we build connections from these to the matrix methods previously covered, while showing how to simplify some of the more arbitrary methods one sometimes learns in a differential equations class.

LINEAR SECOND-ORDER ODEs

Consider the linear 2nd-order autonomous differential equation

$$a \frac{d^2x}{dt^2} + b \frac{dx}{dt} + cx = 0 \quad \Leftrightarrow \quad (aD^2 + bD + cI)x = 0.$$

In a differential equations course, one learns to extract the **characteristic equation**,

$$a\lambda^2 + b\lambda + c = 0,$$

and solve for the characteristic roots, λ_1 and λ_2 . This equation, predicated on the convenient *ansatz* that $e^{\lambda t}$ should be a solution, gives a pair of **basis solutions**, $\phi_{1,2}$, which vary based on the usual trichotomy of distinct/repeated/complex roots:

- ▷ **Real, distinct** $\lambda_1 \neq \lambda_2$: $\phi_1 = e^{\lambda_1 t}$ and $\phi_2 = e^{\lambda_2 t}$.
- ▷ **Real, repeated** $\lambda_1 = \lambda_2 = \lambda$: $\phi_1 = e^{\lambda t}$ and $\phi_2 = te^{\lambda t}$.
- ▷ **Complex** $\lambda_{1,2} = \alpha \pm i\beta$: $\phi_1 = e^{\alpha t} \cos \beta t$ and $\phi_2 = e^{\alpha t} \sin \beta t$.

The general solution is then a linear combination of basis solutions

$$x(t) = C_1\phi_1(t) + C_2\phi_2(t),$$

where the constants, $C_{1,2}$, can be determined from a pair of initial conditions.

This story, often taught in classic differential equations or calculus classes, has a parallel interpretation via linear systems. If we introduce a new variable, y , and set it equal to $y = dx/dt$, then we can convert the linear 2nd-order ODE to a 1st-order system in 2-D:

$$\frac{d}{dt} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} \frac{dx}{dt} \\ \frac{d^2x}{dt^2} \end{pmatrix} = \begin{pmatrix} y \\ -\frac{c}{a}x - \frac{b}{a}y \end{pmatrix} = \begin{bmatrix} 0 & 1 \\ -\frac{c}{a} & -\frac{b}{a} \end{bmatrix} \begin{pmatrix} x \\ y \end{pmatrix}$$

The 2-by-2 matrix, A , above right, has characteristic polynomial $a\lambda^2 + b\lambda + c = 0$. This explains the connection between the approach of factoring the differential operator

$$(aD^2 + bD + cI)x = (D - \lambda_1 I)(D - \lambda_2 I)x = 0,$$

and using the eigenvalues of the matrix A . From our solution to $Dx = Ax$, we have, in the case of real, distinct eigenvalues $\lambda_1 \neq \lambda_2$ with eigenvectors $\mathbf{v}_1, \mathbf{v}_2$,

$$\begin{pmatrix} x \\ y \end{pmatrix} = e^{At} \begin{pmatrix} x_0 \\ y_0 \end{pmatrix} = [\mathbf{v}_1 \quad \mathbf{v}_2] \begin{bmatrix} e^{\lambda_1 t} & 0 \\ 0 & e^{\lambda_2 t} \end{bmatrix} [\mathbf{v}_1 \quad \mathbf{v}_2]^{-1} \begin{pmatrix} x_0 \\ y_0 \end{pmatrix}.$$

As such, $x(t)$ is observed to be a linear combination of the basis solutions $\phi_1 = e^{\lambda_1 t}$ and $\phi_2 = e^{\lambda_2 t}$. Solving for the constant involves knowing the initial conditions. From the work in Chapter 12, one observes the basis solutions for the other (non-real or non-distinct) cases arising in the matrix exponentials.

In general, if given a 2nd-order ODE, it is simplest to work with basis solutions rather than converting everything to matrices, as seen in the examples to follow. However, there will be instances (especially in Volume 3) when an approach using matrices and eigenvalues is advantageous.

SIMPLE HARMONIC OSCILLATOR

It is no surprise that the solutions to the linear oscillator

$$\frac{d^2x}{dt^2} + \omega^2x = 0 \quad \Leftrightarrow \quad (D^2 + \omega^2I)x = 0,$$

consist of periodic waves. This harmonizes with the characteristic roots $\lambda_{1,2} = \pm i\omega$ and the basis solutions $\phi_1 = \cos \omega t$ and $\phi_2 = \sin \omega t$. The general solution is

$$x(t) = C_1 \cos \omega t + C_2 \sin \omega t = C \sin(\omega t - \varphi),$$

where one often converts to a single sine wave of amplitude C and phase φ . The addition of a small amount of friction to the system gives

$$\frac{d^2x}{dt^2} + \nu \frac{dx}{dt} + \omega^2x = 0 \quad \Leftrightarrow \quad (D^2 + \nu D + \omega^2I)x = 0$$

where $\nu > 0$ is a small damping parameter. The characteristic roots are no longer pure imaginary, but have a negative real part:

$$\lambda_{1,2} = \frac{-\nu \pm i\sqrt{4\omega^2 - \nu^2}}{2}.$$

For small amounts of friction $\nu < 2\omega$, the basis solutions are waves with amplitudes that decay like $e^{-\nu/2}$. For $\nu > 2\omega$, the system is *overdamped*, and the solution is an exponential decay to the equilibrium. The case of a repeated root, where $\nu = 2\omega$, corresponds to a *critical damping*.

LINEAR SECOND-ORDER RECURRENCE RELATIONS

The discrete-time version of a 2nd-order equation

$$ax_{n+2} + bx_{n+1} + cx_n = 0 \quad \Leftrightarrow \quad (aE^2 + bE + cI)x = 0,$$

is analogous. The ensuing **characteristic equation**, $a\lambda^2 + b\lambda + c = 0$ has roots λ_1 and λ_2 . The corresponding **basis solutions**, $\phi_{1,2}$, are of the form

- ▷ **Real, distinct** $\lambda_1 \neq \lambda_2$: $\phi_1 = (\lambda_1)^n$ and $\phi_2 = (\lambda_2)^n$.
- ▷ **Real, repeated** $\lambda_1 = \lambda_2 = \lambda$: $\phi_1 = \lambda^n$ and $\phi_2 = n\lambda^n$.
- ▷ **Complex** $\lambda_{1,2} = \alpha \pm i\beta = re^{\pm i\theta}$: $\phi_1 = r^n \cos n\theta$ and $\phi_2 = r^n \sin n\theta$.

The general solution is then a linear combination of basis solutions

$$x_n = C_1\phi_1(n) + C_2\phi_2(n),$$

where the constants, $C_{1,2}$, can be determined from a pair of initial conditions.

If we introduce a new variable, $y_n = x_{n+1}$, then we can convert the linear 2nd-order recurrence relation to a 1st-order system in 2-D:

$$\begin{pmatrix} x_{n+1} \\ y_{n+1} \end{pmatrix} = \begin{pmatrix} x_{n+1} \\ x_{n+2} \end{pmatrix} = \begin{pmatrix} y_n \\ -\frac{c}{a}x_n - \frac{b}{a}y_n \end{pmatrix} = \begin{bmatrix} 0 & 1 \\ -\frac{c}{a} & -\frac{b}{a} \end{bmatrix} \begin{pmatrix} x_n \\ y_n \end{pmatrix}$$

The matrix, A , above, has characteristic polynomial $a\lambda^2 + b\lambda + c = 0$, and the basis solutions arise from the entries in powers of A . For example, in the case of real, distinct eigenvalues $\lambda_1 \neq \lambda_2$ with eigenvectors $\mathbf{v}_1, \mathbf{v}_2$,

$$\begin{pmatrix} x_n \\ y_n \end{pmatrix} = A^n \begin{pmatrix} x_0 \\ y_0 \end{pmatrix} = [\mathbf{v}_1 \quad \mathbf{v}_2] \begin{bmatrix} (\lambda_1)^n & 0 \\ 0 & (\lambda_2)^n \end{bmatrix} [\mathbf{v}_1 \quad \mathbf{v}_2]^{-1} \begin{pmatrix} x_0 \\ y_0 \end{pmatrix}.$$

Other cases follow similarly. As with the continuous-time case, one typically skips conversion to a matrix and jumps straight to the basis solutions. A few examples below illustrate the method and its utility.

FIBONACCI

Consider the classical Fibonacci sequence

$$F = (F_n) = (0, 1, 1, 2, 3, 5, 8, 13, 21, 34, 55, 89, 144, \dots).$$

These numbers – and their connection to the numinous golden ratio – have long been the subject of cosmic speculation. On the more terrestrial plane of this text, the Fibonacci sequence is the solution to the 2nd-order recurrence relation

$$F_n = F_{n-1} + F_{n-2} \quad : \quad F_0 = 0 ; F_1 = 1.$$

To solve this, our first step will be to apply the shift operator, E , twice to obtain

$$F_{n+2} - F_{n+1} - F_n = 0.$$

This can be expressed in operator notation as

$$(E^2 - E - I)F = 0.$$

This factors as

$$(E - \lambda_1 I)(E - \lambda_2 I)F = 0 \quad : \quad \lambda_{1,2} = \frac{1 \pm \sqrt{5}}{2},$$

where $\lambda_{1,2}$ are the roots of the characteristic equation $\lambda^2 - \lambda - 1 = 0$: in this example, the roots are, of course, the golden and silver ratios. The general solution is a linear combination of basis solutions:

$$F_n = C_1 \left(\frac{1 + \sqrt{5}}{2} \right)^n + C_2 \left(\frac{1 - \sqrt{5}}{2} \right)^n.$$

The initial conditions $F_0 = 0, F_1 = 1$ imply that $C_1 = -C_2 = 1/\sqrt{5}$.

Though the solution is complete, it is perhaps a bit unsatisfying to think of the effort of computing the 1000th Fibonacci number using this formula – all those square roots and powers! Fortunately, the lessons of Volume 1 remain: the second eigenvalue, λ_2 , being less than one in absolute value, has powers which rapidly converge to zero. Thus, F_n can be obtained by estimating $C_1(\lambda_1)^n$ using inexpensive logarithms and rounding to the nearest integer.

THE RAISING OF HOGS

The following example of a 2nd-order recurrence relation comes from a market system with time delay. Consider the market for a commodity (let us say hogs) that requires an investment of time before the product comes to market (semiconductors, video games, and other products being analogous). The supply and demand curves for this product are assumed to be affine functions of price P :

$$S(P) = S_0 + aP \quad : \quad D(P) = D_0 - bP$$

In this case, the constants a and b are positive, reflecting the fact that buyers and sellers are price-sensitive: the higher the price of bacon, the more hogs a farmer is motivated to raise.

The subtle point in this model is an (idealized) assumption that it takes exactly two years to raise a hog to market. While the suppliers make their choice based on current prices, the buyers make their choices two years hence. Thus, adding the element of (discrete) time n and equating supply and demand implies the following recurrence relation for equilibrium price:

$$S_0 + aP_n = D_0 - bP_{n+2} \Rightarrow P_{n+2} + \frac{a}{b}P_n = \frac{D_0 - S_0}{b}.$$

This second-order recurrence relation is not linear, but *affine*. There is an equilibrium at

$$P_* = \frac{D_0 - S_0}{a + b}.$$

A change of coordinates to $Q = P - P_*$ converts the affine recurrence relation to the linear relation $(E^2 + a/b I)Q = 0$. Its characteristic equation,


$$\lambda^2 + \frac{a}{b} = 0 \Rightarrow \lambda = \pm i \sqrt{\frac{a}{b}},$$

has pure imaginary roots. Converting to polar form and taking powers gives a general solution

$$Q_n = C_1 \left(\frac{a}{b}\right)^{\frac{n}{2}} \cos \frac{n\pi}{2} + C_2 \left(\frac{a}{b}\right)^{\frac{n}{2}} \sin \frac{n\pi}{2}.$$

This implies a perfect 4-year cycle of prices rising and falling, with suppliers following whatever the current trend is and always over- or under-producing (by an amount dependent on the distance to the equilibrium).

The key is the ratio of supply price sensitivity a to demand price sensitivity b . If



sellers are more price-sensitive than buyers ($a > b$), then $|\lambda| > 1$ and the equilibrium is unstable: suppliers overreact to the latest trend and cyclically overproduce/underproduce, causing price crashes/surges, in a 4-year cycle with things getting worse each cycle. This lag-response cycle is not unique to production processes. One sees this in many time-delayed control systems, such as when you turn the temperature dial in the shower and the flow takes a few seconds to respond. In that case, oversensitivity leads to an unstable spiral of overreaction.

On the other hand, when buyers are more price-sensitive than sellers ($a < b$) then the oscillations damp out over time and the price converges to equilibrium.



CHAPTER 15 : CLASSIFICATION OF EQUILIBRIA

THE INITIAL EMPHASIS on linear-algebraic techniques in linear dynamics obscures the visual elements so helpful in 1-D. This chapter is a segue from the algebraic and analytic to the more geometric and dynamical aspects of linear systems.

ILLUSTRATING PLANAR DYNAMICS

Recall from Chapter 4 when illustrating continuous-time dynamics in 1-D the pivot from graphs of solutions based on time to a more implicit representation based on direction of the dynamics: a *vector field*. In 2-D, things are more complicated, but not unfamiliar. Planar vector fields are familiar from calculus class. They are precisely what a continuous-time 2-D nonlinear system of the form $D\mathbf{x} = F(\mathbf{x})$ specifies. Equilibria of the dynamics are *zeroes*, or locations where the vector field vanishes.

More useful than vector fields are *flows*. Given any point $(x_0, y_0) \in \mathbb{R}^2$, thought



of as an initial condition, the solution to $D\mathbf{x} = F(\mathbf{x})$ is a parametrized curve $(x(t), y(t))$ in the plane, called a *flowline*. These flowlines are oriented (by time) and fill up (or *foliate*) the plane (apart from the equilibria). This echoes 1-D, in which the line $\mathbb{R}^1$ is partitioned by the equilibria into a small set of disjoint open intervals: flowlines. The topology of $\mathbb{R}^2$ is more expansive and allows for greater freedom of flow.

It is a recurring lesson of this Volume that discrete-time systems are harder and less manageable: visualizing 2-D discrete-time dynamics is a challenge.

The continuous-time method is sometimes applicable: *e.g.*, to linear systems of the form $E\mathbf{x} = A\mathbf{x}$ where the eigenvalues of A are both positive. In such a case, one has a foliation of the plane into invariant “flowlines” and the discrete-time dynamics are akin to “jumping” ahead by some time Δt . Negative eigenvalues (that correspond to axis *flips*) and full nonlinear dynamics can break this visualization method. We therefore begin with continuous-time linear systems.


THE GARDEN OF EQUILIBRIA

In 1-D, equilibria are classified by the stable-unstable-degenerate trichotomy, determined by the Stability Criterion applied to the derivative at the equilibrium. Recognizing this derivative as an eigenvalue (for the trivial 1-by-1 matrix derivative) tempts the reader to think that equilibrium types will be

doubled in 2-D. More complexity is warranted, as eigenvalues can come in conjugate pairs.

One classifies the equilibrium at the origin of a linear system qualitatively. These have their own taxonomy, complete with zoological nomenclature.

- ▷ SOURCE : $0 < \lambda_2 < \lambda_1$: unstable along both eigendirections.
- ▷ SINK : $\lambda_2 < \lambda_1 < 0$: stable along both eigendirections.
- ▷ SADDLE : $\lambda_2 < 0 < \lambda_1$: mixed stable-unstable behavior.
- ▷ SPIRAL SOURCE : $\lambda = \alpha \pm i\beta$; $\alpha > 0$: spirals out.
- ▷ SPIRAL SINK : $\lambda = \alpha \pm i\beta$; $\alpha < 0$: spirals in.
- ▷ CENTER : $\lambda = \pm i\beta$; $\beta \neq 0$: foliated by ellipses.



This is not all. There are *degenerate* equilibria, characterized by the presence of a zero eigenvalue. If the other eigenvalue is nonzero, then it introduces either a stable or unstable element to the dynamics. Some authors distinguish fine shades of ever-weaker forms of stability (such as *weak*, *neutral*, or *semi*-); this text admits the existence of degeneracy but avoids dwelling there long enough to product a taxonomy. Also of questionable identification is the case of repeated roots, either stable (negative) or unstable (positive). Though their linear-algebra is burdensome, their qualitative behavior is simple – they are sinks or sources, on the very edge of spiraling into or out of control.

EXAMPLE: Computing the eigenvalues of the following linear system is simple:

$$Dx = Ax \quad : \quad A = \begin{bmatrix} 3 & 0 \\ 8 & -7 \end{bmatrix} \quad \leftrightarrow \quad \frac{dx}{dt} = 3x \quad ; \quad \frac{dy}{dt} = 8x - 7y.$$

They are precisely the diagonal entries, as the matrix is triangular. This continuous-time system thus has a saddle at the origin. The y -axis is the stable eigenvector.

THE TRACE-DETERMINANT METHOD

In 2-D continuous-time linear systems, there is a wonderful method for quickly classifying the equilibrium at the origin. Observe that, for a 2-by-2 matrix A , the

characteristic equation is $\lambda^2 - \text{tr}(A)\lambda + \det A$, where $\text{tr}(A)$ is the **trace** of the matrix. Solving for the eigenvalues gives:

$$\lambda_{1,2} = \frac{1}{2} \left(\text{tr}(A) \pm \sqrt{\text{tr}^2(A) - 4 \det A} \right).$$

The key term is the **discriminant**, $\text{tr}^2(A) - 4 \det A$. When this is negative, then the eigenvalues are complex with nonzero imaginary part. If, in addition, trace vanishes, then the eigenvalues are pure imaginary. A positive discriminant implies real eigenvalues; a zero discriminant implies repeated real eigenvalues.

Combined with the fact that for any square matrix, the determinant is the product of the eigenvalues and the trace is the sum of the eigenvalues, we have a few useful conclusions about the equilibrium at the origin:


- ▷ SADDLE if $\det A < 0$
- ▷ SINK if $\text{tr } A < 0$ and $0 < \det A < (\text{tr } A)^2/4$
- ▷ SPIRAL SINK if $\text{tr } A < 0$ and $\det A > (\text{tr } A)^2/4$
- ▷ SOURCE if $\text{tr } A > 0$ and $0 < \det A < (\text{tr } A)^2/4$
- ▷ SPIRAL SOURCE if $\text{tr } A > 0$ and $\det A > (\text{tr } A)^2/4$
- ▷ CENTER if $\text{tr } A = 0$ and $\det A > 0$
- ▷ DEGENERATE if $\det A = 0$


This leads to a complete determination of the type of equilibrium in 2-D continuous-time linear dynamics based solely on the trace and determinant of the matrix. This is neatly encoded in a **trace-determinant diagram**, to be memorized and used frequently.

EXAMPLE: Computing the eigenvalues of the following linear system would be unpleasant:

$$Dx = Ax \quad : \quad A = \begin{bmatrix} \sqrt{7} & \ln \frac{1}{100} \\ e^5 & -\pi \end{bmatrix}.$$

However, determining the type of equilibrium at the origin is clean. The trace of A is small and negative; the determinant is positive and large relative to trace. This is a spiral sink.





DISCRETE SUBTLETIES

The trace-determinant method of classification is not as well-known in discrete time. The Stability Criterion switches focus from positive/negative to magnitude of eigenvalues. However, the computation of eigenvalues in terms of trace and determinant remains intact: we proceed by partitioning the plane and classifying components. First, the discriminant vanishes along $\text{tr}^2(A) - 4 \det A$, and above this lies complex conjugate eigenvalue pairs. The centers occur in this region where $\det A = 1$; below that lie spiral

sinks; above, spiral sources. The remaining partitions of the plane are governed by the equations

$$\frac{1}{2}(\text{tr}(A) \pm \sqrt{\text{tr}^2(A) - 4 \det A}) = \pm 1,$$

which implies $\det A = -1 \pm \text{tr} A$. These degenerate lines, where one real eigenvalue crosses ± 1 , cuts the trace-determinant plane into saddles, sources, and sinks. The comparison of this diagram with its continuous-time analogue reveals the subtleties of eigenvalues in 2-D.

STABILITY CRITERION REDUX

With the classification of equilibria in linear systems complete, a reformulation of the Stability Criterion is in order. The 2-D is more involved than the 1-D case; however, we can lift the 1-D case to individual eigenvalues, pronouncing them to be stable, unstable, or neutral based on the 1-D Stability Criterion. This leads to the following.

STABILITY CRITERION, CONTINUOUS TIME

For a linear 2-D dynamical system $D\mathbf{x} = A\mathbf{x}$ with eigenvalues $\lambda_{1,2}$, the stability of the equilibrium at the origin is based on the real component of the eigenvalues:

- ▷ STABLE if $\Re(\lambda_i) < 0$ for all i
- ▷ UNSTABLE if $\Re(\lambda_i) > 0$ for some i
- ▷ NEUTRAL if $\Re(\lambda_i) = 0$ for some i

STABILITY CRITERION, DISCRETE TIME

For a linear 2-D dynamical system $E\mathbf{x} = A\mathbf{x}$ with eigenvalues $\lambda_{1,2}$, the stability of the equilibrium at the origin is based on the modulus of the eigenvalues:

- ▷ STABLE if $|\lambda_i| < 1$ for all i
- ▷ UNSTABLE if $|\lambda_i| > 1$ for some i
- ▷ NEUTRAL if $|\lambda_i| = 1$ for some i

It is imperative to examine this carefully. First, the logical subtlety: stability requires all eigenvalues to be stable: any instability destabilizes the system. This is sensible: witness the dynamics near a saddle equilibrium. Next, the distinction between continuous and discrete time is sharpened: what was murky in 1-D becomes clear in the context of the complex plane of potential eigenvalues. The Exponential Lemma of Chapter 2 – that “ $e^D = E$ ” – is the key. The partition of the complex plane into domains of stable, degenerate, and unstable eigenvalues for D , when exponentiated (via *Euler’s formula*), yield the stable-neutral-unstable partition for the shift operator E . This is sublime.

It is also to be noted that there are finer shades of stability and instability one can define. Our use of terms *stable* and *unstable* imply exponential behavior; one might consider subexponential growth or decay as being similarly categorized. In this text, we do not add to the terminology with these emanations of stability (asymptotic, neutral, Lyapunov, *etc.*): stability means linear stability as above.

EXAMPLE: Consider the following linear system in continuous time:

$$D\mathbf{x} = A\mathbf{x} \quad : \quad A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}.$$

The double-zero eigenvalue means that it is a (very) degenerate equilibrium. It has no unstable eigenvalues. Nevertheless, almost all initial conditions “go to infinity”, as can be seen by computing an explicit solution: $\mathbf{x}(t) = C_1 + C_2 t$. This can be regarded as a weak form of instability, but we will not call such an equilibrium [linearly] unstable.

CHAPTER 16 : NONLINEAR 2-D SYSTEMS

LINEAR SYSTEMS have been our primary focus: we are now prepared to turn to fully nonlinear systems in 2-D. The approach will imitate that of 1-D, with a few additional twists; some complications come from the mechanics of linear algebra in 2-D, whereas others come from the less highly constrained topology of the plane.

LINEARIZATION AT EQUILIBRIA

The story from 1-D systems continues to 2-D: given a nonlinear system, we find the equilibria, linearize the dynamics about the equilibria, and then use our knowledge of linear systems, informed by eigenvalues and eigenvectors. Specifically, given

$$D\mathbf{x} = F(\mathbf{x}) \text{ or } E\mathbf{x} = F(\mathbf{x}) \text{ where } \mathbf{x} = \begin{pmatrix} x \\ y \end{pmatrix} \text{ \& } F(\mathbf{x}) = \begin{pmatrix} F_1(x, y) \\ F_2(x, y) \end{pmatrix},$$

one finds an equilibrium $\mathbf{a}$ and linearizes $F(\mathbf{x})$ about $\mathbf{x} = \mathbf{a}$ using the derivative $[DF]$ evaluated at $\mathbf{a}$. As F is a function with two inputs and two outputs, the derivative at $\mathbf{a}$ is a linear transformation represented as a 2-by-2 matrix

$$[DF]_{\mathbf{a}} = \begin{bmatrix} \frac{\partial F_1}{\partial x} & \frac{\partial F_1}{\partial y} \\ \frac{\partial F_2}{\partial x} & \frac{\partial F_2}{\partial y} \end{bmatrix}_{\mathbf{a}}.$$

Some authors insist on a separate name for this (the *Jacobian*), but that is to be avoided: it is simply *the derivative*. Setting $\mathbf{h} = \mathbf{x} - \mathbf{a}$ as a perturbation to an equilibrium, the induced dynamics on this perturbation are given by

$$\begin{aligned} D\mathbf{h} &= D\mathbf{x} - D\mathbf{a} = F(\mathbf{a} + \mathbf{h}) = F(\mathbf{a}) + [DF]_{\mathbf{a}}\mathbf{h} + O(|\mathbf{h}|^2) \approx [DF]_{\mathbf{a}}\mathbf{h} \\ E\mathbf{h} &= E\mathbf{x} - E\mathbf{a} = F(\mathbf{a} + \mathbf{h}) - \mathbf{a} = [DF]_{\mathbf{a}}\mathbf{h} + O(|\mathbf{h}|^2) \approx [DF]_{\mathbf{a}}\mathbf{h} \end{aligned}$$

These are the linearized dynamics about the equilibrium, to which we apply our knowledge of eigenvalues, eigenvectors, and linear solutions. Perturbations to the equilibria grow according to the linear solutions:

$$\mathbf{h}(t) \approx \exp([DF]_{\mathbf{a}}t) \mathbf{h}_0 \quad : \quad \mathbf{h}_n \approx ([DF]_{\mathbf{a}})^n \mathbf{h}_0.$$

Any confusion at this point should be alleviated by rereading Chapter 3 before proceeding.

EXAMPLE: Consider the linear system given by

$$\frac{dx}{dt} = x - y \quad : \quad \frac{dy}{dt} = 1 - xy.$$

There are equilibria where $x = y$ and $xy = 1$; thus at $(1,1)$ and $(-1,-1)$. The derivative of the right-hand side evaluated at these equilibria yields:

$$\begin{bmatrix} 1 & -1 \\ -y & -x \end{bmatrix}_{1,1} = \begin{bmatrix} 1 & -1 \\ -1 & -1 \end{bmatrix} \quad \& \quad \begin{bmatrix} 1 & -1 \\ -y & -x \end{bmatrix}_{-1,-1} = \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix}.$$

In the former case, we have negative determinant and conclude that $(1,1)$ is a saddle. In the latter case, the trace and determinant each equal 2: this implies that $(-1,-1)$ is a spiral source.

SOME ASSEMBLY REQUIRED

In 1-D, continuous-time, knowing what happens at all the equilibria completely determines the rest of the dynamics. This is not the case in 2-D, and it is nontrivial to interpolate between the local dynamics about equilibria.

There are several means of discerning the transition from local to global, including the following:

- ▷ *Pick a Point*: For a spiral or center, knowing the direction of rotation is helpful. Though converting to local polar coordinates and determining the dynamics on the angular coordinate would work, it is usually simpler to pick a single point near the equilibrium and evaluate the dynamics there.
- ▷ *Isoclines*: For continuous-time dynamics, visualizing curves where the vector field takes on a particular slope (*iso-cline*) can be a help. Setting the components $F_1 = 0$ or $F_2 = 0$ (or even $F_1 = \pm F_2$) is often helpful in piecing together the flow of things.
- ▷ *Occam's Razor*: When in doubt, the simplest assembly of local dynamics into a global picture is to be preferred as *ab initio* a good guess. As with other applications of the Razor, simplicity is not guaranteed.

Applying these hints to the previous example indicates that the spiral is spinning counterclockwise and that these spiral trajectories feed into the local dynamics of the saddle. Could there be more happening? Could the linear approximation have been misleading? These are important questions.

THE HARTMAN-GROBMAN THEOREM

The limits of linearization were more apparent in 1-D: degenerate equilibria, at which the derivative yields neither stability nor instability, were clearly not classified by the derivative. In higher dimensions, this story persists, but at the level of eigenvalues: any neutral eigenvalues are sufficient to call linearization into doubt.

One says that an equilibrium is *hyperbolic* if it has no neutral eigenvalues. This terminology, emanating from deep corners of geometry, is central to dynamical systems, and it shall be in continuous use. The result that guarantees the accuracy of linearization utilizes the language of topological conjugacy (see Chapter 5) to compare the qualitative behavior of nonlinear and linearized dynamics.

HARTMAN-GROBMAN THEOREM: On a sufficiently small neighborhood of a hyperbolic fixed point $\mathbf{a}$, the nonlinear dynamics of F (assumed C^1) are topologically conjugate to the dynamics of the linearized system $[DF]_{\mathbf{a}}$.

▷ *Idea*: The proof is nontrivial and involves building a nonlinear change of coordinates to match the nonlinear dynamics with the local linearized dynamics. As one might guess, the Implicit Function Theorem plays a starring role. The continuous-time and discrete-time proofs are very similar in spirit. ◁

This means in practice that one can be confident in the presence of sources, sinks, spirals, and saddles, but centers are a suspicious occurrence. *Nota bene*: this is a *local* result which applies only to a neighborhood of a hyperbolic equilibrium.

STABLE & UNSTABLE CURVES

With a bit more attention to detail in the application of the Implicit Function Theorem, one can extract some additional information from the proof of the Hartman-Grobman Theorem. In the case of a hyperbolic equilibrium of saddle-type, the stable and unstable eigenspaces of the linearized dynamics are one-dimensional and point in the directions along which the dynamics are stable and unstable respectively. What about the nonlinear dynamics? The stable and

unstable eigenspaces are linearizations of *stable* and *unstable curves* for the dynamics. For a saddle equilibrium $\mathbf{a}$ define the following:

- ▷ STABLE CURVE: $W^s(\mathbf{a}) = \{x \in \mathbb{R}^2 : x \rightarrow \mathbf{a} \text{ as time goes to } +\infty\}$.
- ▷ UNSTABLE CURVE: $W^u(\mathbf{a}) = \{x \in \mathbb{R}^2 : x \rightarrow \mathbf{a} \text{ as time goes to } -\infty\}$.

The notation may seem odd: it will be generalized greatly and put to general use in Volumes 3 and 4. It is a deep result that these stable and unstable curves are indeed curves, not merely locally but globally. The *Stable Manifold Theorem* of Volume 3 will cover this and more.

INTEGRALS & CENTERS

When applying the Hartman-Grobman Theorem, it certainly *feels* correct that one should not trust degenerate equilibria: there's simply not enough information. Centers – in contrast – give the appearance of being well-behaved. Why are these not to be trusted? Do they ever truly exist?

EXAMPLE: Consider the following continuous time system:

$$D \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 2y + 4y^3 \\ -2x - 4x^3 \end{pmatrix}.$$

This system has an equilibrium at the origin whose linearization indicates a center. This is in fact a true nonlinear center, as is hinted by a simulation. To confirm, one makes the inspired choice to investigate the function $\Phi = x^2 + y^2 + x^4 + y^4$. The level sets of Φ coincide with the orbits of the flow since, by the Chain Rule,

$$\frac{d\Phi}{dt} = \frac{\partial \Phi}{\partial x} \frac{dx}{dt} + \frac{\partial \Phi}{\partial y} \frac{dy}{dt} = (2x + 4x^3)(2y + 4y^3) - (2x + 4x^3)(2y + 4y^3) = 0.$$

This means that flowlines maintain a constant value of Φ and, indeed, sweep out level sets. Since Φ has a minimum at the origin, it is surrounded by simple closed curves: a center.

A system is said to be *integrable* (or *conservative*) if there is a (locally non-constant) *integral* $\Phi: \mathbb{R}^2 \rightarrow \mathbb{R}$ such that Φ does not change along orbits of the dynamics. Assuming (as we shall) that Φ is continuously differentiable, this means that

- ▷ CONTINUOUS: $D\Phi = 0$; that is, $\frac{d}{dt} \Phi(x(t)) = 0$.
- ▷ DISCRETE: $E\Phi = 0$; that is, $\Phi(x_{n+1}) = \Phi(x_n)$.

The existence of an integral greatly constrains the types of behavior seen in 2-D.

LEMMA: Orbits of integrable 2-D systems are contained within level sets of the integral Φ . Equilibria of such systems can be saddles, centers, or degenerate: no sinks or sources – regular or spiral – exist.

▷ *Proof*: The first statement follows from the definition of the integral. The second statement can be argued by contradiction: in a neighborhood of any sink or source, such a Φ must by continuity and invariance be a constant, violating the non-constant assumption. ◁

The problem with this method – as you may have guessed – is that there is no oracle to tell you what the integral Φ might be if you suspect that one exists.

There is one method for finding a conserved quantity which resonates with certain ideas from multivariable calculus. A 2-D continuous-time system is said to be *Hamiltonian* if there is a (C^1) function $H: \mathbb{R}^2 \rightarrow \mathbb{R}$ such that

$$\frac{dx}{dt} = \frac{\partial H}{\partial y} \quad : \quad \frac{dy}{dt} = -\frac{\partial H}{\partial x}.$$

Hamiltonian systems in 2-D are integrable (with integral H), since mixed partials commute. The converse is not true: there are many integrable systems which are not Hamiltonian. However, on the plane, Hamiltonian-ness is detectible.

LEMMA: A continuous-time (C^1) system $D\mathbf{x} = F(\mathbf{x})$ on $\mathbb{R}^2$ is Hamiltonian with (a C^2) integral H if and only if its divergence is zero: $\nabla \cdot F = 0$.

▷ *Idea*: Necessity follows from the definition of Hamiltonian and the fact that mixed partial derivatives commute. Sufficiency comes from a 2-D version of the *Poincaré Lemma* from multivariable calculus (any divergence-free field is a curl).

◁

CHAPTER 17 : POPULATION MODELS

NONLINEAR SYSTEMS can exhibit a great many types of equilibria. Our immediate goal is to put this expressiveness to work in modelling complicated phenomena. This chapter focuses on population models.

Most of the models in this chapter are attributed to Lotka and Volterra; the term *Lotka-Volterra model* is common, but often ambiguous, as there are several different types of population models. We therefore use more descriptive titles.

PREDATOR-PREY MODEL

The following models two species. The first, x , grows naturally in its environment and is modelled with linear dynamics $Dx = r_1x$ for $r_1 > 0$ a constant. The second species, y , is predative, and, in the absence of food, dies off; thus, $Dy = -r_2y$ for $r_2 > 0$. Such an uncoupled linear system clearly gives a saddle equilibrium at the origin, with unbounded growth in x and eventual extinction for y .

To make the model more interesting, assume that y is predator and x prey. Whenever predators and prey meet, there is a benefit to y and a detriment to x . The rate of predator-prey encounters depends on the relative sparsity of populations; very little prey or very few predators means that an encounter will be less frequent. The model therefore assumes that encounters occur at a rate proportional to xy , with constants of proportionality tuning both frequency and net benefit/detriment. The *predator-prey model* emerges:

$$\begin{aligned}\frac{dx}{dt} &= x - \alpha xy \\ \frac{dy}{dt} &= -ry + \beta xy\end{aligned}$$

Here, we have rescaled populations so that the prey has unit reproduction rate and predators have a relative ratio $r > 0$, with the minus sign forcing natural decline. The predator-prey detriment/benefit rates are (resp.) $\alpha, \beta > 0$.

Let us follow the story. There is an equilibrium at $(0,0)$, and, since the linear portion is clear, one observes that it is a saddle with eigenvalues $\lambda_1 = 1$ and $\lambda_2 = -r$. Note as well that the x and y axes are the (invariant) stable and unstable curves – once you start there, you never leave. This is a typical occurrence in population models, since populations do not burst forth *ex nihilo*.

There is a second equilibrium at $x = r/\beta$ and $y = 1/\alpha$. The derivative of the right hand side of the model at this point is

$$\begin{bmatrix} 1 - \alpha y & -\alpha x \\ \beta y & \beta x - r \end{bmatrix}_{\substack{x=r/\beta \\ y=1/\alpha}} = \begin{bmatrix} 0 & -\frac{r\alpha}{\beta} \\ \beta & 0 \end{bmatrix}.$$

With trace zero and determinant positive, we are met with a center as the linearized dynamics. Our next step – to invoke the Hartman-Grobman Theorem and declare victory – falls flat, as this is one of the fragile cases in which we cannot say for sure that the linearization is accurate, as higher-order terms not seen by the linearization can destroy a center. Does the Hamiltonian method of Chapter 15 imply that it is a center? No, it does not. Nevertheless, this is a true nonlinear center.

LEMMA: The predator-prey model is integrable with integral

$$\Phi = \beta x - r \ln x + \alpha y - \ln y.$$

▷ *Proof*: Direct computation shows that $d\Phi/dt = 0$.

$$\frac{d\Phi}{dt} = \beta x(1 - \alpha y) - \frac{rx(1 - \alpha y)}{x} + \alpha y(\beta x - r) - \frac{y(\beta x - r)}{y} = 0. \quad \triangleleft$$

It is not hard to see that level sets of Φ in the first quadrant are simple closed curves surrounding the equilibrium as illustrated. The entire first quadrant is filled and foliated with periodic orbits.

What this means for the population model is that, away from equilibrium, populations of predators and prey rise and fall periodically, though not in sync. As with the time-delay hogs model from Chapter 13, an undersupply of predators leads to an oversupply of prey, which encourages greater numbers of predators, *etc.* One curious feature of the model is that for certain initial conditions of too many predators and prey, the population sizes take turns collapsing to nearly zero before slowly building back to same inflated numbers.

COMPETITIVE MODEL

Instead of an asymmetric predator-prey model, one can work with a model in which the two species, x and y , compete for shared resources. The following 2-species competitive model has, like the previous model, a single parameter, $r > 0$, encoding a relative reproduction rate, as well as parameters $\alpha, \beta > 0$ which record the relative impact of competitive encounters between species. Populations are normalized so that 1 is the carrying capacity of each population in the absence of the other. Beginning with decoupled logistic models and adding in negative quadratic interaction terms as before yields the competitive (Lotka-Volterra) model:

$$\begin{aligned} \frac{dx}{dt} &= x(1 - x - \alpha y) \\ \frac{dy}{dt} &= ry(1 - y - \beta x) \end{aligned}$$

This system requires a more substantive analysis. There are four possible equilibria, located at $(0,0)$, $(1,0)$, $(0,1)$, and

$$\left(\frac{1 - \alpha}{1 - \alpha\beta}, \frac{1 - \beta}{1 - \alpha\beta} \right),$$

which is a “realistic” equilibrium only when α and β are both less than one or both greater than one. The derivative of the dynamics is the following matrix:

$$\begin{bmatrix} 1 - 2x - \alpha y & -\alpha x \\ -r\beta y & r - 2ry - r\beta x \end{bmatrix}$$

The four equilibria are classified as follows.

1: $(0,0)$ is always a source, since the derivative here is

$$\begin{bmatrix} 1 & 0 \\ 0 & r \end{bmatrix},$$

and $r > 0$. This means that when population sizes are low, there is no competition: both species are free to reproduce at their natural rates.

2: $(1,0)$ is a sink for $\beta > 1$ and a saddle for $\beta < 1$, since the derivative is

$$\begin{bmatrix} -1 & -\alpha \\ 0 & r(1 - \beta) \end{bmatrix}.$$

and $r > 0$. Since β measures the inhibitory impact of the x species on y , larger β is beneficial to establishing x as the dominant species.

3: $(0,1)$ is a sink for $\alpha > 1$ and a saddle for $\alpha < 1$: a symmetric analysis to $(1,0)$.

4: The final equilibrium, when both populations have nonzero size, is of most interest. As an exercise, you can show:

- ▷ This is a sink when $\alpha, \beta < 1$.
- ▷ This is a saddle when $\alpha, \beta > 1$.

When $\alpha, \beta < 1$, we have a state of *cooperation* or *coexistence*: neither species is stable in its full size, but both admit sustainable “market shares”. When $\alpha, \beta > 1$, we have a critical situation called *competitive exclusion*: there can be only one winner. With work, including computation of the stable and unstable curves to the saddle, one gets additional information. The unstable curve limits at the two winner-take-all population sinks. The stable curve forms a *separatrix* – a boundary that separates the two futures.

Note that if you find yourself near the stable curve to a saddle in the context of competitive exclusion, you may be lulled into supposing that you are in a period of cooperation, coexistence, or stalemate. It may take a while for the unstable eigenvalue to kick in. When it does, you will be dismayed at how fast your fortunes can change, whether in business, biology, or war.

DISCRETE-TIME VERSIONS

Replacing the continuous-time differentiation operator with the discrete-time forward difference yields the following version of the competitive model:

$$\begin{aligned}x_{n+1} - x_n &= x_n(1 - x_n - \alpha y_n) \\y_{n+1} - y_n &= r y_n(1 - y_n - \beta x_n)\end{aligned}$$

The analysis is similar: there are again four equilibria at the exact same locations. The difference is that we have extra terms on the diagonal of the derivative, due to moving the x_n and y_n over to the right hand side. The derivative is:

$$\begin{bmatrix} 2 - 2x - \alpha y & -\alpha x \\ -r\beta y & 1 + r - 2ry - r\beta x \end{bmatrix}$$

The rest of the classification is similar; the trace-determinant method is still helpful here.

It is an exercise for the reader to repeat this discretization with the predator-prey model. One should find that the analysis of equilibria is similar. What is different is in the periodic orbits. In discrete-time, there is still a center, but the invariant loops can exhibit rational or irrational rotation behavior, depending on initial conditions.

POPULATION COHORT MODELS

Instead of using the freedom of two dimensions to model a pair of species in an adversarial or symbiotic relationship, one can work with two cohorts within a population. Such can occur in linear or nonlinear dynamics in discrete or continuous time.

One famous model for infectious disease is the *SIR model*, which decomposes a population of size N into three cohorts: susceptible, infected, and recovered, of respective sizes $S + I + R = N$. The resulting system of dynamics is (after eliminating the R variable) given by

$$\frac{d}{dt} \begin{pmatrix} S \\ I \end{pmatrix} = \begin{pmatrix} -\beta IS/N \\ \beta IS/N - \gamma I \end{pmatrix},$$

where $\beta, \gamma > 0$ are transmission and remission parameters. Their ratio, $R_0 = \beta/\gamma$, is the all-too-familiar **reproduction number** of the infectious disease.

Here, the limitations of 2-D are apparent. One frequently needs more than two cohorts in order to properly model populations. For example, in epidemiology, one might care about additional cohorts E (exposed), C (asymptomatic carrier), and D (deceased). Chapter 26 in Volume 3 will go into greater detail.

CHAPTER 18 : NONLINEAR OSCILLATORS

MUCH OF THE INITIAL WORK in 20th century dynamics came from engineers and physicists seeking to understand nonlinear oscillators. This chapter sets up several classical models, expressed as continuous-time second-order nonlinear differential equations. Though this will give an opportunity to practice classification of equilibria, the limits of our existing tools will soon be seen, prompting a turn to additional theoretical work.

NONLINEAR PENDULUM

The simple harmonic oscillator of Chapter 13 is a woefully inadequate model, good only for small angles. A more accurate model of a rigid-rod pendulum whose bob makes an angle θ with the vertical axis is

$$mL^2 \frac{d^2\theta}{dt^2} = -mgL \sin \theta .$$

As a first-order system, this becomes:

$$\begin{pmatrix} d\theta/dt \\ dv/dt \end{pmatrix} = \begin{pmatrix} v \\ -\frac{g}{L} \sin \theta \end{pmatrix}$$

There are equilibria where $v = 0$ and $\sin \theta = 0$. Since $\theta \in \mathbb{S}^1$ is an angular variable (recall Chapter 5), we have precisely two equilibria: $(0,0)$ and $(0, \pm\pi)$, with the system best viewed as having state space equal to a cylinder $\mathbb{S}^1 \times \mathbb{R}^1$. The derivative of the right-hand-side is simple to compute:

$$\begin{bmatrix} 0 & 1 \\ -\frac{g}{L} \cos \theta & 0 \end{bmatrix} .$$

Given that the trace is zero, the classification of equilibria depends only on θ . Linearization declares the equilibrium at $\theta = 0$ to be a center; the equilibrium at $\theta = \pm\pi$ is a saddle. The former is suspicious, but ultimately accurate. The lack of friction points to an integral for the (Hamiltonian) system given by

$$H = \frac{1}{2} v^2 - \frac{g}{L} \cos \theta .$$

Orbits of this system thus lie on level sets of H . As the origin is a local minimum of H , it is a true center.

The second equilibrium is worth investigation. This equilibrium corresponds to a perfectly still vertical pose: clearly unstable. However, the saddle point $\mathbf{p}$ does have stable and unstable curves (cf. Chapter 15). A perfectly executed kick from the bottom can send this pendulum converging to the vertical as time slowly unfolds. With slightly less energy, the pendulum will forever swing back and forth in large arcs. With more, the pendulum winds about the circular state *ad infinitum*.

The exceptional behavior on display can be summarized as:

$$W^s(\mathbf{p}) \cap W^u(\mathbf{p}) \neq \emptyset .$$

In other words, the stable and unstable curves to the vertical pose intersect in an orbit that limits to $\mathbf{p}$. Such an orbit which converges in both forwards and backwards time to a saddle equilibrium is an example of a **homoclinic orbit**: it *inclines* to the *same* equilibrium solution in directions of time.

SPINNING PENDULUM

A slight generalization leads to some interesting complications. Take this same frictionless planar pendulum and spin the rig about a vertical axis at a constant

angular velocity, ω , here treated as a parameter. The relevant equation of motion (with mass cancelled out) is:

$$L \frac{d^2\theta}{dt^2} = L\omega^2 \sin \theta \cos \theta - g \sin \theta .$$

As a first-order system, this becomes:

$$\frac{d\theta}{dt} = v \quad : \quad \frac{dv}{dt} = \omega^2 \sin \theta \cos \theta - \frac{g}{L} \sin \theta .$$

The system has equilibria at $(0,0)$, $(\pm\pi, 0)$, and, assuming the angular speed ω to be sufficiently high, $(\arccos g/L\omega^2, 0)$. Note that since $\theta \in \mathbb{S}^1$, there is a single equilibrium at $\theta = \pm\pi$. In addition, one must be careful about the inverse cosine function: there are two values of $\arccos g/L\omega^2$. To what physical configurations do these correspond?

Given the previous example, one would guess that the equilibrium at $\theta = \pm\pi$ is a saddle, and that at $\theta = 0$ is a center. To verify, one computes the derivative as

$$\begin{bmatrix} 0 & 1 \\ \omega^2(\cos^2 \theta - \sin^2 \theta) - \frac{g}{L} \cos \theta & 0 \end{bmatrix} .$$

The trace is always zero, so linearization gives either saddles or centers depending on the sign of the determinant. For $\omega^2 > g/L$, both these equilibria are saddles: sufficiently fast rotation destabilizes the origin, sending nearby initial conditions on a swing. The two remaining equilibria, where the bob is still and rotating with the frame, are each centers. The lack of friction in this problem portends that the system has an integral, and that these are true centers.

VAN DER POL OSCILLATOR

Simple electrical circuits give rise to systems of linear differential equations. The introduction of a non-linear circuit element can prompt a novel twist in even a simple layout. Such is the case for the Van der Pol circuit, in which a nonlinear resistor in an RLC circuit yields a simple, unforced *Van der Pol oscillator*:

$$\frac{d^2x}{dt^2} + \epsilon(x^2 - 1) \frac{dx}{dt} + x = 0 .$$

Here, $\epsilon > 0$ is a small constant. This resembles a simple harmonic oscillator with a damping term that is dependent on the size of x : there is positive damping (friction) for large values of x and negative damping (excitation) for small values. Converting this to a first-order system yields:

$$\begin{pmatrix} \dot{x} \\ \dot{y} \end{pmatrix} = \begin{bmatrix} 0 & 1 \\ -1 & \epsilon \end{bmatrix} \begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} 0 \\ -\epsilon x^2 y \end{pmatrix} .$$

For $\epsilon > 0$ small there is a unique equilibrium at the origin. With a small positive trace, and a positive determinant, one concludes a spiral source. This would seem to be all that can be said. The existence of positive damping far off would merit examining a larger neighborhood of the origin. Upon so doing, one finds the existence of an attracting periodic orbit – a *limit cycle* – is apparent in simulations. Proving that such an orbit exists is a nontrivial task: see Chapter 20.

DUFFING OSCILLATOR

The Duffing oscillator is of great classical interest and is again a modification of a damped simple harmonic oscillator. To the linear term a cubic term is added, yielding:

$$\frac{d^2x}{dt^2} + \frac{dx}{dt} + \delta x + x^3 = 0.$$

The physically interesting case occurs where $\delta < 0$. Physical examples to which this model applies include various *double well potential* mechanical systems, such as an elastic beam being pulled by a pair of magnets or springs. Converting this to a first-order system yields:

$$\begin{pmatrix} \dot{x} \\ \dot{y} \end{pmatrix} = \begin{bmatrix} 0 & 1 \\ -\delta & -1 \end{bmatrix} \begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} 0 \\ -x^3 \end{pmatrix}.$$

There is clearly an equilibrium at the origin and, with the linear part clear, the classification is seen to depend upon the parameter δ . For $\delta < 0$, the origin is a saddle; for $\delta > 0$, the origin is stable (since trace is negative). It is a sink for $0 > \delta \geq 1/4$, and a spiral sink for $\delta > 1/4$. The other equilibria of this system occur where $y = 0$ and $x(\delta + x^2) = 0$; thus at $x = \pm\sqrt{-\delta}$. This pair of equilibria only exist for $\delta < 0$; the derivative evaluated at these equilibria equals:

$$\begin{bmatrix} 0 & 1 \\ -\delta - 3x^2 & -1 \end{bmatrix}_{\pm\sqrt{-\delta}, 0} = \begin{bmatrix} 0 & 1 \\ 2\delta & -1 \end{bmatrix}.$$

This has trace -1 and determinant $-2\delta > 0$: sinks, spiraling for $\delta < -\frac{1}{8}$.

PERIODICALLY FORCED

Many of the oscillators in this chapter have especially interesting behavior in the nonautonomous setting, where the right-hand side of the ODE is time-dependent. For example, a *periodically forced Duffing oscillator* is:

$$\frac{d^2x}{dt^2} + \frac{dx}{dt} + \delta x + x^3 = A \cos \omega t,$$

where A is the *amplitude* of the forcing and ω its *frequency*. Such systems are beyond what we can as yet handle: though second-order, this system's time-varying dynamics increments the dimension. It is especially interesting to visualize solutions evolving in time. Though the unforced case has a pair of sinks (for $\delta < 0$), the forced oscillator does not converge to an equilibrium state. This will be one motivation for our work in Volume 4.

CHAPTER 19 : BIFURCATIONS REDUX

PARAMETERS are essential components of most good models; witness the population and oscillator models of the previous two chapters. As a result, one expects to see bifurcations – qualitative changes in the dynamics as a function of parameters. The next step after classification of equilibria is the classification of bifurcations of equilibria.

DECOUPLED SYSTEMS & BIFURCATIONS

To a large degree, bifurcations in 2-D systems recapitulate the 1-D classification. Our strategy for classifying equilibria in 2-D systems has involved performing a coordinate change to decouple the linearized system. The corresponding strategy for classifying bifurcations works similarly, though the actual coordinate changes are, as in 1-D bifurcation theory, more subtle than this text can manage at full detail.

The simplest way to lift the bifurcation types from Volume 1 to 2-D is to embed the bifurcation in the following decoupled system,

$$D \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} f(x, \mu) \\ \lambda y \end{pmatrix} \quad : \quad E \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} f(x, \mu) \\ \lambda y \end{pmatrix},$$

where $f(x, \mu)$ expresses the normal form of the bifurcation with parameter μ . Here, $\lambda \neq 0$ is the eigenvalue for the (linear) dynamics on the second variable.

For example, the continuous-time 2-D saddle-node bifurcation can be expressed in a decoupled form as:

$$\frac{dx}{dt} = \mu - x^2 \quad : \quad \frac{dy}{dt} = \lambda y.$$

This system expresses the usual saddle-node in the x variable, having either no equilibria ($\mu < 0$) or two equilibria of opposite stabilities ($\mu > 0$). Together with the linear dynamics in y , the system has a pair of equilibria at $(\sqrt{\mu}, 0)$ which collide at $\mu = 0$ and disappear as μ becomes negative.

What types of equilibria are these? This depends on λ . For $\lambda > 0$, the two equilibria which appear are a source and a saddle, but for $\lambda < 0$, the equilibria are a saddle and a sink respectively. This explains at last the etymology of the term *saddle-node*: a **node** is an outdated term for a source or a sink. In 2-D, the saddle-node creates one saddle and one node.

In general, one does not (want to) perform a coordinate change to decouple the system. The investigation of equilibria suffices to identify the relevant bifurcation.

EXAMPLE: The following parametrized system is a better version of the genetic switch considered in Chapter 8. Here, $x(t)$ is related to the amount of a protein produced and $y(t)$ is related to messenger RNA, with the pair evolving in continuous time as

$$\begin{aligned} \frac{dx}{dt} &= y - ax \\ \frac{dy}{dt} &= -by + \frac{x^2}{1 + x^2} \end{aligned}$$

The parameters $a, b > 0$ are both positive. There is always an equilibrium at the origin; the linearized dynamics has eigenvalues $\lambda_1 = -a$, $\lambda_2 = -b$. There is also a pair of equilibria at

$$\left(\frac{1 \pm \sqrt{1 - 4a^2b^2}}{2ab}, \frac{1 \pm \sqrt{1 - 4a^2b^2}}{2b} \right).$$

This pair of equilibria exists only when the term in the square root is nonnegative: $2ab \leq 1$. Moving across the hyperbola $ab = 1/2$ in the parameter plane changes the system from one equilibrium to three. Is this a pitchfork bifurcation? No. The pair of new equilibria appear at $(1/2ab, 1/2b)$, not the origin. Thus, this bifurcation is a curve of saddle-nodes.

OSCILLATIONS REDUX

Several interesting bifurcations occur in models of nonlinear oscillators, seen in Chapter 17. These can be a little complicated in the case of an integrable system, but the standard classification of 1-D bifurcations is usually sufficient.

EXAMPLE: Recall the Duffing oscillator, written as a first-order system:

$$D \begin{pmatrix} x \\ y \end{pmatrix} = \begin{bmatrix} 0 & 1 \\ -\delta & -1 \end{bmatrix} \begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} 0 \\ -x^3 \end{pmatrix}.$$

For $\delta < 0$, the origin is a saddle; it is a sink for $0 < \delta \leq 1/4$, and a spiral sink for $\delta > 1/4$. The other equilibria appear at $(\pm\sqrt{-\delta}, 0)$ for $\delta < 0$. These are sinks, changing to spiral sinks for $\delta < -1/8$. The interesting bifurcation happens at the origin at $\delta = 0$, where the two sinks collide with the saddle at the origin, converting it to a sink. This is a pitchfork bifurcation. The preponderance of stability implies supercriticality.

EXAMPLE: Recall the spinning planar pendulum:

$$\frac{d\theta}{dt} = v \quad : \quad \frac{dv}{dt} = \omega^2 \sin \theta \cos \theta - \frac{g}{L} \sin \theta.$$

The system has equilibria at $(0,0)$, $(\pm\pi, 0)$, and $(\arccos g/L\omega^2, 0)$. There are two values of $\arccos g/L\omega^2$ for $\omega^2 > g/L$, and each of the associated equilibria are centers. At the critical angular speed where $\omega^2 = g/L$, these two centers collapse to the origin, and as ω decreases, the origin changes from a saddle to a center. This is a supercritical pitchfork bifurcation at the origin when the parameter ω is at $\sqrt{g/L}$. Though similar to the Duffing oscillator, this is different: the lack of friction gives an integral which forces equilibria to be either saddles or centers.

POPULATIONS REDUX

Transcritical bifurcations are very common in population models, due to the biological constraint that a zero-size population is always an equilibrium.

EXAMPLE: Recall the continuous-time competitive model of Lotka-Volterra from Chapter 16. The four equilibria were classified: $(0,0)$ is always a source, but the other three,

$$(1,0); (0,1); \left(\frac{1-\alpha}{1-\alpha\beta}, \frac{1-\beta}{1-\alpha\beta} \right)$$

were either sinks or saddles depending on the values of α and β . It is a worthwhile exercise to argue that this system experiences a transcritical bifurcation at $(1,0)$ when $\alpha = 1$ and at $(0,1)$ when $\beta = 1$. This, then, explains what happens to the fourth equilibrium when it fails to be physically realistic: it has passed through one of the transcritical bifurcations.

Other types of bifurcations are of course possible, including types that we have not yet encountered.

EXAMPLE: The predator-prey model of Lotka-Volterra from Chapter 16 has a discrete-time version (due to J. M. Smith):

$$x_{n+1} = rx_n(1 - x_n) - x_n y_n \quad : \quad y_{n+1} = x_n y_n / c.$$

Here, $c, r > 0$ are parameters. This system has equilibria at $(0,0)$ and at $(c, r(1-c) - 1)$, meaning that one wants $r(1-c) > 1$. This nonzero equilibrium has derivative

$$\begin{bmatrix} r - 2rx - y & -x \\ y/c & x/c \end{bmatrix}_{c, r(1-c)-1} = \begin{bmatrix} 1 - cr & -\beta \\ \frac{1}{c}(r - rc - 1) & 1 \end{bmatrix}.$$

This has trace $2 - cr$ and determinant equal to $r(1 - 2c)$; this system undergoes a discrete-time bifurcation when $cr < 4$ and $r = 1/(1 - 2c)$, changing between spiral sink and spiral source. This is a fascinating and an instructive system from the point of view of bifurcations and periodic orbits, as will be seen in the next chapter.

FORESHADOWING

As indicated by the previous example, our story of bifurcations is by no means complete. Everything thus far has centered on equilibria and how one unfolds changes in the number and types of equilibria about a singularity. However, there is one exception to this – a foreshadowing of things to come.

In discrete time 1-D systems, the *period-doubling* bifurcation implicated not only an equilibrium, but a periodic orbit as well: the periodic orbit sprang forth from the equilibrium at a change of stability. This required a negative derivative, or, more suggestively, a local *flipping* of the reals about the equilibrium. That local rotation by half-a-turn in turn drove a novel bifurcation implicating period-2 orbits.

In 2-D systems, in discrete or continuous time, there is more room for flipping and rotation in general, and room must be made for the bifurcations this enables. No longer constrained by the topology of the line to suppress periodic orbits, planar dynamics in continuous time will evince bifurcations with and from periodic orbits, as will be seen in the next chapter and the remainder of this Volume.

CHAPTER 20 : THE HOPF BIFURCATION

BIFURCATIONS can implicate more than merely equilibria. One of the genuinely novel features of 2-D dynamics arises when considering bifurcations associated to systems that are “oscillatory” in nature. For most this chapter, we restrict to the case of continuous-time systems, noting at the end the vast increase in complexity when adapting to discrete-time.

THE NORMAL FORM

All the bifurcations thus far presented in the Volume are uncoupled, in the sense that, after a change of coordinates, there is a 1-D system with bifurcation that is independent of what is happening in the other dimension. This implies that all the equilibria involved have real eigenvalues (since, otherwise, independent directions cannot be decoupled).

Consider the following parametrized system whose linearization at the origin has complex eigenvalues:

$$\begin{pmatrix} \frac{dx}{dt} \\ \frac{dy}{dt} \end{pmatrix} = \begin{bmatrix} \mu & -\omega \\ \omega & \mu \end{bmatrix} \begin{pmatrix} x \\ y \end{pmatrix} + c \begin{pmatrix} x(x^2 + y^2) \\ y(x^2 + y^2) \end{pmatrix} \quad (*)$$

Here $\omega \neq 0$ and $c \neq 0$ are constants and μ is a parameter. This system has a unique equilibrium at the origin. Its linearization is explicit from the equation: the eigenvalues can be read off as $\lambda_{1,2} = \mu \pm i\omega$. When $\mu < 0$, there is a spiral sink; it becomes a spiral source when $\mu > 0$.

What happens at $\mu = 0$ is a bifurcation, since the type of the equilibrium changes. Linearization predicts a center, but the Hartman-Grobman Theorem counsels caution: the nonlinear terms might break the center when $\mu = 0$.

THE BIRTH OF A LIMIT CYCLE

To see what happens, it is best to convert to polar coordinates. This is motivated by the appearance of $r^2 = x^2 + y^2$ in the nonlinear terms. Recall the formulæ for converting Euclidean derivatives to polar:

$$\frac{dr}{dt} = \frac{1}{r} \left(x \frac{dx}{dt} + y \frac{dy}{dt} \right) \quad : \quad \frac{d\theta}{dt} = \frac{1}{r^2} \left(x \frac{dy}{dt} - y \frac{dx}{dt} \right).$$

Substituting in (*) yields

$$\frac{dr}{dt} = \mu r + cr^3 \quad : \quad \frac{d\theta}{dt} = \omega.$$

This is a decoupled system. The angular coordinate evolves like a spinner, rotating at speed $\omega \neq 0$. It is the radial coordinate that is interesting. There is an equilibrium at $r = 0$ which is stable for $\mu < 0$ and unstable for $\mu > 0$, corresponding to the change from spiral sink to spiral source. Noting that $r \geq 0$ in polar coordinates, we observe that there is an equilibrium in the radial variable:

$$r_* = \sqrt{-\mu/c} \quad : \quad \mu/c < 0.$$

For the appropriate values of μ (opposite the sign of c), there is a radial equilibrium $r_* > 0$. If you begin at an initial condition (x_0, y_0) in the plane with $x_0^2 + y_0^2 = r_*^2$, what happens? The radial coordinate is invariant, while the angle θ evolves at a constant nonzero rate: you trace out a circle in the Euclidean plane. This periodic orbit is different than those seen in a center: they are limit cycles. A **limit cycle** is a periodic orbit to which nearby initial conditions are attracted in

either forward or backward time. To see that this happens with a Hopf bifurcation, note that the stability of the radial equilibrium at $r = r_*$ is given by linearization at r_* :

$$\frac{d}{dr}(\mu r + cr^3)|_{r_*} = \mu + 3cr_*^2 = \mu - 3\mu = -2\mu.$$

Thus, for $\mu > 0$, the limit cycle is *stable* and attracts nearby orbits; for $\mu < 0$, the limit cycle is *unstable* and repels solutions away from it.

What happens in practice is that the periodic orbit emerges suddenly and grows rapidly in radius. This is explained via the formula $r_* = \sqrt{-\mu/c}$: the rate of change of the radius with respect to the parameter at the bifurcation is infinite. One experiences this sudden change viscerally in physical instances of the bifurcation.

SUPERCritical VS. SUBCRITICAL

Note that the Hopf bifurcation depends critically on the sign of coefficient, c , of the cubic terms. The resemblance to a pitchfork bifurcation is more than coincidence: Hopf bifurcations come in supercritical and subcritical variants.

A **supercritical Hopf**, where $c < 0$, is characterized by a spiral sink becoming unstable and throwing off a stable limit cycle. A **subcritical Hopf**, where $c > 0$, has the spiral source stabilizing and throwing off an unstable limit cycle. In either case, one can distinguish super/subcriticality by examining what happens at $\mu = 0$, whether it is a weakly spiraling sink or source.

The difference between the two is not academic. Subcritical Hopf bifurcations, like their pitchfork counterparts, are genuinely dangerous. You may have a system in equilibrium at a spiral sink – a small perturbation oscillates with exponentially decreasing amplitude. If you are very close to a Hopf bifurcation and change your parameter just past the bifurcation, what happens? The stable spiral sink becomes an unstable spiral source, and the oscillations you think should damp out in fact begin to grow in amplitude.

Your future now depends critically on which Hopf you have. If this was a *supercritical* Hopf bifurcation, the amplitude of your oscillations levels out as you converge to a stable limit cycle. Once you notice this, you can in principle *turn the dial back* and reverse the bifurcation, with your oscillations now damping back to the stable equilibrium. However, if your Hopf was *subcritical*, there is no nearby limit cycle on this side of the bifurcation: your oscillations will continue to increase in amplitude. Perhaps, noting this, you decide to *turn the dial back* to where it was before the bifurcation. *Too late!* Your new initial condition is no longer very close to the equilibrium. If you are outside the unstable limit cycle, then your oscillations will continue to grow, even though your parameter is back to where it was when you were safe.

THE UBIQUITOUS HOPF

The Hopf bifurcation is of codimension one, and thus is generic in 1-parameter families of systems. Of all the bifurcations covered in this text, the continuous-time Hopf bifurcation is likely the most widely observed. Consider the following types of behaviors, all of which can arise at a critical transition between equilibrium and periodic states.

- ▷ *Vibrating*: When driving an old car on the highway, it can happen that near a certain speed, the vehicle begins to vibrate and rumble.

- ▷ *Shuddering*: A similar phenomenon happens with a faulty shopping cart wheel, wherein at a critical speed, it begins shaking in its direction.
- ▷ *Fluttering*: Wind blowing through blinds covering an open window can induce a fluttering of the blinds that is both visible and audible.
- ▷ *Fishtailing*: Towing a trailer can be dangerous, depending on the weight distribution. Perturbations to bearing can either damp out or swing wildly out of control.
- ▷ *Dripping*: Very carefully controlling the flow rate of a faucet appears to give a bifurcation between steady, laminar flow, and regular dripping.
- ▷ *Knocking*: In old pipes, changing the flow rate of a faucet can lead to a loud knocking, with regular cavitation of the flow within. One can hear the amplitude of the resulting limit cycle.
- ▷ *Rolling*: Atmospheric fluid with a weak temperature gradient (warm on the bottom; cool on the top) conducts heat smoothly. With a sharp enough gradient, the warm fluid below rises, inducing convection rolls and periodic flow.
- ▷ *Shedding*: Flow over an airplane wing at low speeds is smooth and laminar. As the speed increases, it can happen that vortices are generated periodically.
- ▷ *Shaking*: Musculo-skeletal control systems that you use to steady your hands can, in certain circumstances, fail in such a way as to give tremors. In patients with Parkinson's disease, this can be endemic and debilitating.
- ▷ *Firing*: Certain models of electrochemical potentials in neurons exhibit behaviors that range from steady-state to periodic firing based on parameters.
- ▷ *Dying*: Depending on growth rates and resource consumption rates of predator-prey models, populations can either equilibrate or fall into boom-bust cycles of resource depletion, famine, death, and regrowth.

Which of these truly express Hopf bifurcations? This depends on the precise details of the continuous-time models used to describe them. Without models, examples can quickly become speculative to an immoderate degree. Are bipolar mood swings the result of a Hopf bifurcation, and with respect to what parameter? Is the business cycle the result of a Hopf bifurcation? Without giving in too much to unscientific speculation, knowing how to identify bifurcations in systems in which you have no explicit model is a worthwhile skill to develop.

A CHEMICAL OSCILLATOR

Perhaps you recall the canonical chemistry experiment [the Briggs-Rauscher reaction] where the right combination of reactants (typically hydrogen peroxide and iodate) change colors in a spectacular oscillatory manner: that would suggest a Hopf bifurcation. One simple model for a chemical oscillator is the following:

$$\frac{dx}{dt} = 1 - (b + 1)x + ax^2y \quad : \quad \frac{dy}{dt} = bx - ax^2y$$

The variables $x, y \geq 0$ represent concentrations, and there are two positive parameters $a, b > 0$. This system has a unique equilibrium located at $x = 1$ and $y = b/a$. The linearized dynamics has derivative:

$$\begin{bmatrix} -b - 1 + 2axy & ax^2 \\ b - 2axy & -ax^2 \end{bmatrix}_{1, b/a} = \begin{bmatrix} b - 1 & a \\ -b & -a \end{bmatrix}.$$

We expect a Hopf bifurcation to take place when the determinant is positive and the trace vanishes: this is precisely where $b = a + 1$. Consequently, we have a line in the (a, b) plane along which a Hopf bifurcation is expected. When $b < a + 1$, the trace is negative and we have a spiral sink; for $b > a + 1$, a spiral source. Experiments would suggest that for $b > a + 1$ there is a stable limit cycle – the chemical oscillator. However, we cannot conclude that rigorously without knowing more.

A CRITERION

Determining super- or sub-criticality is not easy. Fortunately, there is an analytic criterion that, if not straightforward, is at least simple to implement. If, at the parameter value at which a Hopf bifurcation takes place, the system in local coordinates at the equilibrium is of the following form:

$$D \begin{pmatrix} x \\ y \end{pmatrix} = \begin{bmatrix} 0 & -\omega \\ \omega & 0 \end{bmatrix} \begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} f(x, y) \\ g(x, y) \end{pmatrix},$$

where f and g are in $O(x^2 + y^2)$ (that is, they are of order quadratic and higher terms), then to determine the nature of the Hopf bifurcation, compute the following combination of partial derivatives of f and g :

$$f_{xxx} + f_{xyy} + g_{xxy} + g_{yyy} \\ + \frac{1}{\omega} (f_{xy}(f_{xx} + f_{yy}) - g_{xy}(g_{xx} + g_{yy}) - f_{xx}g_{yy} + f_{yy}g_{xx}).$$

LEMMA: If the quantity above is positive, the bifurcation is subcritical; if negative, it is supercritical.

▷ *Idea:* The proof is an unpleasant exercise in manipulating Taylor series, polar coordinates, and time, in order to extract the coefficient of the r^3 term. ◁

EXAMPLE: Consider the following second order ODE

$$\frac{d^2x}{dt^2} + \mu \frac{dx}{dt} + x - x^2 \frac{dx}{dt} + x^3 = 0.$$

Writing this as a first order system with $y = dx/dt$ yields

$$\begin{pmatrix} x \\ y \end{pmatrix} = \begin{bmatrix} 0 & 1 \\ -1 & -\mu \end{bmatrix} \begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} 0 \\ x^3 - x^2y \end{pmatrix}.$$

With $\omega = -1$, $f = 0$, and $g = x^3 - x^2y$, one computes from the criterion that this equation has a subcritical Hopf bifurcation at the origin.

IN DISCRETE TIME

It is perhaps an unusual departure for this text to ignore discrete time when handling a subject so important as that of a Hopf bifurcation. Is it because there is no discrete-time analogue? No. In a 2-D discrete time system, when an equilibrium transitions from spiral sink to spiral source, there is a bifurcation that is very like the Hopf. It is different enough to merit its own name: the **Neimark-Sacker bifurcation**. It is difficult enough to merit its own place at the end of the Volume, where most readers will not mind a truncated overview.

The simplest example of a Neimark-Sacker bifurcation at the origin is best written out directly in polar coordinates:

$$r_{n+1} = (1 + \mu)r_n + cr_n^3 \quad : \quad \theta_{n+1} = \theta_n + \omega,$$

where $c \neq 0$ is real and $\omega \in \mathbb{S}^1$ are constants. From this decoupled system, so long as ω is not *poorly chosen*, one has a spiral sink at the origin for $\mu < 0$ small and a spiral source for $\mu > 0$ small. The pitchfork-like bifurcation in the (nonnegative) radial direction comes in supercritical ($c < 0$) and subcritical ($c > 0$) varieties. In either case, a family of invariant attracting (super-) or repelling (sub-) circles of radius $\sqrt{-\mu/c}$ springs from the origin at the bifurcation.

EXAMPLE: Recall the predator-prey model from Chapter 18:

$$x_{n+1} = rx_n(1 - x_n) - x_ny_n \quad : \quad y_{n+1} = x_ny_n/c.$$

The equilibrium at $(c, r(1 - c) - 1)$ has derivative

$$\begin{bmatrix} r - 2rx - y & -x \\ y/c & x/c \end{bmatrix}_{c, r(1-c)-1} = \begin{bmatrix} 1 - cr & -\beta \\ \frac{1}{c}(r - rc - 1) & 1 \end{bmatrix}.$$

When $cr < 4$, $c < 4/9$, and $r = 1/(1 - 2c)$, there is a Neimark-Sacker bifurcation. Numerical simulation indicates that it is supercritical.

Though this would appear to be a simple discrete-time analogue of the Hopf, there are several points at which the narrative diverges from that of the cleaner continuous-time story: Chapter 22 will take up the story.

CHAPTER 21 : FINDING PERIODIC ORBITS

PERIODIC ORBITS and limit cycles are just as important as equilibria in understanding global qualitative features of dynamical systems. Unfortunately, they are by no means easy to detect. This chapter focuses on tools for establishing or eliminating the existence of periodic orbits and limit cycles in planar systems.

THE BENDIXSON-DULAC CRITERION

When trying to find a periodic orbit, it is sometimes best to check whether such is even possible. The following criterion is one approach to ruling out existence.

THEOREM: Let $\vec{F} = F(x, y) = f(x, y)\hat{i} + g(x, y)\hat{j}$ be a continuously-differentiable vector field on a simply-connected domain $U \subset \mathbb{R}^2$. If, for some scalar field $\rho: U \rightarrow \mathbb{R}$, the divergence $\nabla \cdot (\rho\vec{F})$ is nonzero on U , then the continuous-time dynamics induced by $\vec{F}$ has no periodic orbits contained in U .

▷ *Proof:* This is a classic application of Green's Theorem. Assume that γ is a periodic orbit of $\vec{F}$ in U ; then $\vec{F}$ (and $\rho\vec{F}$) is everywhere tangent to γ . The (topological) disc D of which γ is the boundary lies within U by simple-connectivity. The flux of $\vec{F}$ (and $\rho\vec{F}$) out of D is zero by tangency, and Green's Theorem implies that

$$0 = \int_{\gamma} \rho\vec{F} \cdot \vec{n} \, d\ell = \iint_D \nabla \cdot (\rho\vec{F}) \, dA \neq 0.$$

This is a contradiction to the assumption. ◁

As with most results based on Green's Theorem, when it works it is magical, but it rarely works. Frustration comes from the infinite supply of possible rescalings ρ , with no algorithm for how to choose.

EXAMPLE: The simple version of Bendixson-Dulac does not work with the dynamics

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

since the divergence is $-1 + 2y$. However, rescaling the vector field by $\rho = e^{-2x}$ works, as the divergence is now $-2e^{-2x}y + e^{-2x}(-1 + 2y) = -e^{-2x} < 0$. There are no periodic orbits.

Where did this $\rho = e^{-2x}$ come from? That is the crux of the difficulty.

THE POINCARÉ-BENDIXSON THEOREM

Ruling out the existence of periodic orbits is not always possible, especially when they exist. In 2-D continuous, time, there is a remarkably powerful theorem for detecting limit cycles. This result uses the topology of the plane to constrain the possible limiting behaviors of forward-time orbits which do not escape to infinity.

POINCARÉ-BENDIXSON THEOREM: A forward-time orbit of $D\mathbf{x} = F(\mathbf{x})$ which lies within a closed bounded region of $\mathbb{R}^2$ either (1) is a periodic orbit; (2) limits to a periodic orbit; or (3) limits to an equilibrium.

[The phrase *limits to* has a technical sense: details do not in these margins fit.]

▷ *Idea:* There is no simple proof of Poincaré-Bendixson: it is a deep result relying on the topology of the plane. The key technical steps are the construction of a Poincaré return map (see Volume 3) to a recurrent orbit, combined with a fixed point theorem that implies a limit cycle. ◁

The interested reader is welcome to discover the elements of the proof by taking a page of paper (the bounded region) and drawing a flowline which winds about, never crossing itself nor slowing down.

This theorem is especially useful for forcing the existence of limit cycles. One common application involves the setting of a trap to lure orbits to a periodic prison.

EXAMPLE: The following system has a difficult-to-find periodic orbit:

$$\frac{dx}{dt} = 2x - 2y - 2x^3 - 3xy^2 \quad : \quad \frac{dy}{dt} = 2x + 2y - x^2y - y^3.$$

There is a unique equilibrium at the origin. Converting to polar coordinates yields

$$\begin{aligned} \frac{dr}{dt} &= \frac{1}{r}((2x^2 - 2xy - 2x^4 - 3x^2y^2) + (2xy + 2y^2 - x^2y^2 - y^4)) \\ &= \frac{1}{r}(2r^2 - 2x^4 - 4x^2y^2 - y^4) = 2r\left(1 - r^2 - \frac{1}{2}r^2 \sin^4 \theta\right). \end{aligned}$$

Consider the annular region $\frac{1}{2} \leq r \leq 2$ in the plane. At $r = \frac{1}{2}$, $\frac{dr}{dt} > 0$, and when $r = 2$, $\frac{dr}{dt} < 0$. This region is **forward-invariant**: any point on the boundary is pushed within and cannot escape. As there are no equilibria in this annulus, there must, by Poincaré-Bendixson, be a limit cycle within this annulus. Checking the monotonicity of $d\theta/dt$ shows that this in fact encircles the origin in a counter-clockwise fashion.

Such an inescapable equilibrium-free maze for orbits to run is known as a **trapping region**. One can (with more delicate and difficult work) find a trapping region for the Van der Pol oscillator of Chapter 17 and confirm the existence of the limit cycle hinted at.

GRADIENT vs HAMILTONIAN DYNAMICS

Certain classes of dynamics can be more tractable with respect to periodic orbits, exhibiting either an abundance or scarcity in turn.

A dynamical system in continuous or discrete time is said to be a **gradient system** if the right hand side is exactly a gradient for some $H: \mathbb{R}^2 \rightarrow \mathbb{R}$:

$$Dx = -\nabla H(x) \quad : \quad Ex = -\nabla H(x).$$

The negative sign is artificial but suggestive of the lazy nature of a downhill sink into minima. In 2-D (and beyond) gradient systems are devoid of periodic orbits, due to the interpretation of the gradient vector as the direction of maximal rate of increase of a function.

The situation is different for Hamiltonian systems. Recall from Chapter 13 the special case of a Hamiltonian system in continuous time, which we may write using matrix-vector notation as the following:

$$Dx = -J\nabla H(x) \quad : \quad J = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}.$$

The matrix J is the representation of the imaginary unit in real 2-by-2 matrices. It has the effect of twisting a gradient system into one that runs along level sets of H instead of across them. Such an H is thus, as per Chapter 13, an *integral* of the system. There is a preponderance of periodic orbits, though no limit cycles, as the dynamics are area-preserving.

GLOBAL BIFURCATIONS & LIMIT CYCLES

Our treatment of bifurcation theory in Chapter 18 continued the story from Volume 1 by focusing on bifurcations in a neighborhood of a singular equilibrium. Such **local bifurcations** are the foil of the more generally wild **global bifurcations** which often implicate periodic orbits. A complete classification of such is impossible; however, a few foundational examples are worth contemplation.

[HB] HOMOCLINIC BIFURCATION. Recall from the example of the nonlinear pendulum in Chapter 17 that a **homoclinic orbit** is one which limits to an equilibrium, $\mathbf{a}$, in both directions of time, forwards and backwards. Consider a continuous-time system with a homoclinic orbit to a saddle. This is an exceptional circumstance, as the stable and unstable curves of the saddle must globally evolve in such a manner as to perfectly line up and overlap. A small perturbation to the system would cause the stable and unstable curves to slightly miss.

There is a codimension-one global bifurcation which breaks a homoclinic orbit into a limit cycle. The parameter μ corresponds (roughly) to the amount of separation between $W^s(\mathbf{a})$ and $W^u(\mathbf{a})$. On one side of the bifurcation lies a limit cycle nearing the saddle $\mathbf{a}$, whose period increases without bound. At the critical parameter, the stable and unstable curves match, and the limit cycle becomes a homoclinic orbit – an “infinite period” cycle. What happens as the parameter passes through the bifurcation value is that the stable and unstable curves separate in the other direction: the limit cycle is no more.

[SNP] SADDLE-NODE of PERIODIC ORBITS. Consider a pair of nested limit cycles in the plane – one stable and one unstable. These can experience a global bifurcation of saddle-node type in which, as the parameter is varied, the two cycles come closer and eventually merge and annihilate. A simple decoupled example can be written out in polar coordinates:

$$\frac{dr}{dt} = r(\mu - (r - C)^2) \quad : \quad \frac{d\theta}{dt} = \omega.$$

Here, ω and $C > 0$ are constants and μ is a parameter. One can easily spot the limit cycles as solutions to $dr/dt = 0$. In this case, there are a pair of circular limit cycles at radii $C \pm \sqrt{\mu}$ for $\mu > 0$. These collide and annihilate at $r = C$ when μ vanishes.

Both of the above are examples of global bifurcations: one cannot localize them to take place within the neighborhood of an equilibrium. Although the homoclinic bifurcation implicates an equilibrium, from a local perspective, nothing has changed during the bifurcation. In each case it is the global interaction with a limit cycle that matters.

BOGDANOV-TAKENS BIFURCATION

A brief exposition of one interesting codimension-2 bifurcation ties together much of what we have seen. Consider the 2-D continuous-time system

$$\frac{dx}{dt} = y \quad : \quad \frac{dy}{dt} = \alpha + \beta y + x^2 + xy,$$

where α and β are parameters. The bifurcation unfolds about the origin and where both parameters vanish. This is sometimes called a **double-zero-eigenvalue** bifurcation for reasons that should be clear. The reader should as an exercise find and classify the equilibria in this system, with the goal of showing the following:

- ▷ For $\alpha > 0$, there are no equilibria.
- ▷ The line $\alpha = 0$ is a curve of saddle-node bifurcations.
- ▷ There is a curve of Hopf bifurcations where $\beta = \sqrt{-\alpha}$.

What is less straightforward is to determine where the limit cycles exist. Does Bendixson-Dulac help here? What happens to the limit cycles created in the Hopf? With (euphemistically) a *little more work* it can be shown that:

- ▷ The Hopf bifurcations are supercritical, and a stable limit cycle exists for β just below the curve $\sqrt{-\alpha}$ in the (α, β) plane.
- ▷ There is a curve of homoclinic bifurcations along $\alpha \approx -\frac{29}{45}\beta^2 + \sqrt{\beta^3}$.

This curve of homoclinic bifurcations confines the limit cycles to a narrow wedge of parameter space. If you pick parameters at random, you are likely to miss it.

CHAPTER 22 : TOWARDS INDEX THEORY

EXISTENCE IS rarely straightforward. Proving that limit cycles exist or do not exist can be a challenge, and even equilibria can be evasive. In 2-D continuous-time systems, there is an elegant self-contained theory that serves as a gateway to more advanced topological methods in dynamical systems. This chapter will give a brief introduction to *index theory*.

A 1-D PRECURSOR

Index-theoretic methods are topological in nature. This means that they depend upon continuity properties and qualitative features: minimal assumptions. What is the price of a method that gives strong existence results from such simple inputs with nearly no computational effort? The price – as usual – is an increase in abstraction.

As a warm-up example, consider a system of the form $Dx = f(x)$ for $x \in \mathbb{R}$ and f continuous. If one knows the value of f at a pair of points, say, $a < b$, then what can be said? Since this system has no periodic orbits, the best one can hope for is information about equilibria. Assume for simplicity that there are a finite number of equilibria. An application of Rolle's Theorem and a little thought should suffice to convince the reader of the following:

- ▷ If $f(a) < 0 < f(b)$, then there is an unstable equilibrium within (a, b) .
- ▷ If $f(a) > 0 > f(b)$, then there is a stable equilibrium within (a, b) .
- ▷ If $f(a)$ and $f(b)$ have the same sign, nothing can be concluded. There may or may not be any equilibria in the interval (a, b) .

What is interesting is how this lifts to 2-D. Instead of sampling a vector field at a pair of points (the boundary of an interval in the line), one or more simple closed curves (the boundary of a domain in the plane) is needed. The reader who is reminded of Green's Theorem is prepared for what comes next. The reader who is not should perhaps take a moment and review that topic from the calculus of vector fields.

INDEX OF A LOOP

Index in 2-D continuous-time systems is a perfect subject – just difficult enough to be interesting, but close enough to ideas from vector calculus as to not be obtuse. In what follows, we use the term *loop* to mean a simple closed curve that is piecewise differentiable. Throughout, we assume a fixed 2-D continuous-time dynamical system on $\mathbb{R}^2$,

$$\frac{dx}{dt} = f(x, y) \quad : \quad \frac{dy}{dt} = g(x, y),$$

where f and g are continuously-differentiable (C^1) functions of x and y . Given any loop γ that avoids equilibria (that is, f and g do not simultaneously vanish on any point in the image of γ), the following path integral (using the language of 1-forms familiar from vector calculus) yields a well-defined integer called the *index of γ* :

$$I_\gamma = \frac{1}{2\pi} \oint_\gamma \frac{f dg - g df}{f^2 + g^2}.$$

LEMMA: The index of γ is an integer.

▷ *Proof:* The claim, along with a valuable interpretation of I_γ , comes from the following observation. Think of the dynamical system as giving a vector field on the plane, $\vec{F} = f(x, y)\hat{i} + g(x, y)\hat{j}$. Let $\phi(x, y)$ denote the angle made by the vector

$f(x, y)\hat{i} + g(x, y)\hat{j}$ (as measured counterclockwise from the positive x direction). The differential of ϕ , $d\phi$, can then be computed as

$$d\phi = d\left(\arctan \frac{g}{f}\right) = \frac{f dg - g df}{f^2 + g^2}.$$

Thus, the index of γ is expressible as a **winding number**,

$$I_\gamma = \frac{1}{2\pi} \oint_\gamma d\phi,$$

that measures how many full rotations the vector field expresses as you walk along the loop γ . As written, the path integral measures the total change in angle of the vector field along γ . Since γ is a loop, this is an integer multiple of 2π . $\triangleleft$

EXAMPLES OF INDEX

Given the geometric interpretation of index as a total change in angle, the best way to do computations is visually. A few examples lead to several observations:

- ▷ The index of a small loop that does not encircle any equilibria is 0.
- ▷ The index of a CCW loop that encircles a source or a sink is +1.
- ▷ The index of a CCW loop that encircles a saddle is -1.
- ▷ The index of a CCW loop that is also a periodic orbit is +1.
- ▷ Depending on the dynamics, any integer is attainable as the index of a loop.

The first example above generalizes from a local observation to a global proposition.

LEMMA: The index of any loop γ that encircles no equilibria is zero.

▷ *Proof:* Since γ encircles no equilibria, the domain $D \subset \mathbb{R}^2$ that γ bounds satisfies the hypotheses of Green's Theorem; namely, that the integrand of the index,

$$d\phi = \frac{f dg - g df}{f^2 + g^2}$$

is well-defined on D (the denominator does not vanish). The reader familiar with differential forms notation need merely compute that $d(d\phi) = 0$. More explicitly, Green's Theorem, along with a change-of-variables, implies

$$I_\gamma = \frac{1}{2\pi} \oint_\gamma \frac{f}{f^2 + g^2} dg - \frac{g}{f^2 + g^2} df = \iint_D \frac{\partial}{\partial f} \left(\frac{f}{f^2 + g^2} \right) + \frac{\partial}{\partial g} \left(\frac{g}{f^2 + g^2} \right) dA = 0.$$

This completes the proof. $\triangleleft$

INDEX OF AN EQUILIBRIUM

The proof of the above lemma is more broadly useful in establishing the index as a *topological invariant* of the equilibrium: it is independent of the loop chosen.

LEMMA: The index of an isolated equilibrium p is well-defined and independent of the (CCW oriented, simple) loop used to encircle p .

▷ *Proof:* Let γ be a given simple closed curve encircling p (and no other equilibria). Choose a sufficiently small circle u centered at p . Let $\tilde{\gamma}$ denote a piecewise-smooth loop which executes the following:

1. It starts at a point a on γ closest to p ;
2. It follows γ in a counterclockwise direction back to a ;
3. It then traverses a straight path c from a to a point b on u ;
4. follows u in a clockwise orientation back to b ;

5. traverses the path c from b back to the starting point a .

This loop $\tilde{\gamma}$ bounds a region which encloses no equilibria. By the previous lemma, its index is zero. Additivity of path integrals yields:

$$0 = I_{\tilde{\gamma}} = \frac{1}{2\pi} \oint_{\tilde{\gamma}} d\phi = \frac{1}{2\pi} \left(\oint_{\gamma} d\phi + \int_c d\phi - \oint_u d\phi - \int_c d\phi \right) = I_{\gamma} - I_u.$$

Any two loops surrounding only p thus have the same index. $\triangleleft$

As a result, we denote by $I(p)$ the *index* of the (isolated) equilibrium.

ADDITIVITY OF INDEX

Having seen the efficacy of Green's Theorem in this instance, one wants to put that wondrous theorem to work. Only a little additional effort is required to achieve a significant computational result: an additivity theorem.

THEOREM: If γ bounds a disc D containing a finite number of equilibria, its index is computed as

$$I_{\gamma} = \sum_{p \in D} I(p).$$

$\triangleright$ *Proof:* As there are a finite number of equilibria, each is isolated. Modify γ to a loop $\tilde{\gamma}$ that consists of γ interspersed with straight-line excursions to tiny circles about the (finite, hence isolated) equilibria as per the proof of the previous Lemma. The index of $\tilde{\gamma}$ is zero (it bounds no equilibria), and the linear excursions cancel to give the difference between I_{γ} and the sum of the indices over the interior equilibria. $\triangleleft$

This has many useful applications.

COROLLARY: Any periodic orbit encloses at least one equilibrium whose index is positive.

$\triangleright$ *Proof:* By the additivity of index. Note that the sum of indices over an empty set is automatically zero. $\triangleleft$

The reader is encouraged to try and draw an example of a periodic orbit that surrounds only a saddle equilibrium: this will reveal much about the index, the Poincaré-Bendixson Theorem, and more.

INDEX AT INFINITY

The previous proposition – the additivity of index – is part of a deep topological theorem that connects many areas of Mathematics. The first hints of this come when defining a notion of index at infinity.

Several of the ideas and techniques of this chapter are reminiscent of those in an elementary complex analysis course – witness the similarity to contour integrals and the *Cauchy Residue Theorem*. By analogy with completion of the complex plane $\mathbb{C}$ to the Riemann sphere $\mathbb{S}^2$, one might contemplate what happens when *completing* a vector field from the plane to the sphere by adding an abstract equilibrium at “infinity”. What is the index at infinity? This can be computed explicitly and inelegantly via a change of coordinates that amounts to a reflection of a radial coordinate.

The following theorem allows one to work directly with vector fields on a sphere (working in local planar patches to define the index) or to easily compute the index at infinity in terms of all other local indices.

POINCARÉ-HOPF THEOREM: For any continuous vector field on a sphere $\mathbb{S}^2$ with a finite number of equilibria, the sum of the indices of all equilibria equals two.

Thus, for a field on $\mathbb{R}^2$, the fixed point at infinity has index equal to two minus the sum of all other indices of equilibria. The reason that the net sum is precisely two stems from the *Euler characteristic* of the sphere. This takes us to the boundary of algebraic topology, a subject too vast and wonderful for this thin text.

CHAPTER 23 : 2-D MYSTERIES

PLANAR 2-DIMENSIONAL DYNAMICS is the perfect setting for learning: it is complex enough to be able to model and capture interesting phenomena (saddles, spirals, periodic orbits, Hopf bifurcations...) while being simple enough to admit straightforward analysis (trace-determinant methods, Poincaré-Bendixson, index theory...). The apparent simplicity has been carefully arranged; just past the borders of the maps laid out lie intimidating if not dangerous entities. This chapter gives a quick peek at what lies beyond. The conclusion of Volume 1 (Chapter 9) contained a warning not to mistake chaotic pictures and emergent numbers for the deep truths. The need for such cautions has not diminished (and will recur prominently in Volume 4), but this Volume ends on a more adventurous note, as we are presently within reach of several very deep ideas that begin in Dynamical Systems and resonate throughout Mathematics.

BIFURCATION THEORY BEYOND HOPF

Our story of bifurcations has broadened somewhat in this Volume from the simple set of four (SN/TC/PF/PD) 1-D bifurcations to include a few new items. The ubiquitous Hopf bifurcation is the newest, clearest example: as a codimension-1 bifurcation in 2-D continuous-time systems, it is even more commonly identifiable than most 1-D bifurcations. This is, however, not a fully 2-dimensional bifurcation. Recall that when written in polar coordinates, the Hopf is revealed to be a pitchfork bifurcation on the radial variable (with the angular variable uninvolved). This explains why Hopfs come in both supercritical and subcritical variants.

A fully 2-D bifurcation is the discrete-time version of the Hopf known as the *Neimark-Sacker* bifurcation. This also is codimension-1 and at first resembles the Hopf in most respects: in the supercritical case, a spiral sink equilibrium changes to a spiral source, ejecting an invariant closed curve in the plane that surrounds the equilibrium. That closed curve – the discrete time analogue of a stable limit cycle – is attracting but not at all a single orbit. There are typically multiple periodic orbits of a fixed period, but there may be no periodic orbit at all (recall the irrational rotations of Chapter 5). The periods change with parameter, meaning that there are many (many!) subsidiary bifurcations of periodic orbits implicated. The simple example given at the end of Chapter 19 belies the complexity of the Neimark-Sacker bifurcation, whose general form (in polar coordinates) is:

$$r_{n+1} = (1 + \mu + A(\mu)r_n^2)r_n + B(\mu, r_n)r_n^4 \quad : \quad \theta_{n+1} = \theta_n + \omega(\mu) + C(\mu, r_n)r_n^2,$$

where $\omega \in \mathbb{S}^1$ and A, B, C are smooth functions. Unlike in the continuous-time Hopf bifurcation, the parametric dependence of higher-order coefficients has an impact. There are a great many subtle phenomena present, including *resonances* and *Arnol'd tongues*, which themselves are entwined with number-theoretic properties of the rotation number on the invariant curve. For such *Diophantine conditions* and other related numerical intricacies, the curious reader will need to consult other references.

FROM LOCAL TO GLOBAL

This Volume has also witnessed our first examples of *non-local* (that is, *global*) bifurcations, not discernable through a simple Taylor expansion. These include saddle-node bifurcations of periodic orbits [SNP] and homoclinic bifurcations [HB], the latter featuring prominently in the Bogdanov-Takens bifurcation as a new method of birthing (or killing) a periodic orbit. Both of these examples were given in continuous-time systems.

The other simple 1-D bifurcations also arise as global bifurcations of periodic orbits. For example, in a transcritical bifurcation of periodic orbits [TCP], a pair of periodic orbits converge, collide, and “trade stabilities” with respect to one direction. Pitchfork bifurcations of periodic orbits can also occur. Is there a period-doubling bifurcation of periodic orbits? Yes, but only for systems (in continuous time) of dimension three and higher, the reason being a deep connection to the famous *Möbius strip* and its unembeddable nature in 2-D.

Are there other global bifurcations? Certainly. The problem is that one cannot use Taylor expansion and singularity theory as a guide to how to proceed inductively, cataloguing all that can occur. Instead, one must be led by a combination of imagination and numerical investigation in the search for global bifurcations. Such duck-hunts can be very difficult, as in the case of a *canard*, an elusive short-lived limit cycle in certain singularly-perturbed problems. For example, in a singular version of the Van der Pol oscillator of Chapter 17,

$$\frac{dx}{dt} = \frac{1}{\epsilon} \left(y - \frac{x^3}{3} + x \right) \quad : \quad \frac{dy}{dt} = \mu - x ,$$

there are stable limit cycles for $0 < \epsilon \ll 1$ and $\mu \approx 0.998740451245 \dots$ which persist only in minute slivers of the parameter space, to the point where they are nearly impossible to detect computationally, appearing and disappearing in a effervescence of global bifurcations.

Global bifurcations in the context of discrete time are, even in 2-D, too terrible to detail. Volume 4 will trace the chaos that can ensue in the context of a simple-seeming homoclinic bifurcation in 2-D discrete time.

INDEX THEORY BEYOND GREEN

Several times in this Volume, we have seen methods that work very well for continuous-time dynamics – trace-determinant classification, the Poincaré-Bendixson Theorem, the Bendixson-Dulac criterion – but for which no discrete-time analogue exists. Such would seem to be the case for index theory as well. Note the sensitive dependence on Green’s Theorem for all the index-theoretic results proved in Chapter 21.

Is index theory unique to 2-D continuous time? In discrete time on the plane, a great many things break down. One can have a periodic orbit of any period (greater than one) without there being an equilibrium in the system. If there is an invariant closed curve – a loop that bounds a disc – then there must indeed be an equilibrium contained within. However, such an equilibrium may be a nondegenerate saddle, unlike the case of continuous time.

The surprising truth is that index theory makes sense in all dimensions in both discrete as well as continuous time. The setting of 2-D continuous time is the most beautiful and easily intuited case, but index theory more broadly interpreted is a deep and powerful subfield of Mathematics.

There are a number of indices for equilibria in discrete-time systems in 2-D and beyond: the *Morse index*, the *Lefschetz index*, and the *Conley index* are three of the most classical and powerful in the context of dynamics. All these indices require tools from *algebraic topology* which this text cannot properly invoke.

THE GHOSTS OF CHAOS

Both this Volume and the next are explicitly postponing the careful examination of chaotic dynamics that Volume 1 suggested. In 2-D continuous-time systems,

the Poincaré-Bendixson Theorem guarantees that chaotic dynamics (whatever that might entail) are forbidden. Such guarantees are impossible in discrete-time systems. Even slightly nonlinear 2-D discrete-time systems on the plane, such as the well-known Hénon system,

$$\begin{pmatrix} x_{n+1} \\ y_{n+1} \end{pmatrix} = \begin{bmatrix} 0 & 1 \\ 0.3 & 0 \end{bmatrix} \begin{pmatrix} x_n \\ y_n \end{pmatrix} + \begin{pmatrix} -1.4 x_n^2 \\ 0 \end{pmatrix},$$

can possess dynamics which are complex to a degree that makes analysis bewildering: *chaos*. This and other chaotic systems in 2-D discrete time will be a major focus in Volume 4.

VOLUME 3

MANY DIMENSIONS



CHAPTER 24 : GENERAL LINEAR SYSTEMS

DIMENSIONAL CONSTRAINTS have been thus far the critical bound on the dynamics we have considered. The rigidities of line and plane measure out complexities of behavior in coffee-spoon portions. In this Volume, we abandon such constraints to allow for arbitrary dimensions. The flood of contingent phenomena will force a retreat to nearly linear systems, to be more slowly and carefully transcended in next and final Volume.

OUR STORY CONTINUES

The lift from linear and planar dynamics to dimension-independence prompts a recourse to more serious linear algebra, culminating in the *Jordan canonical form*. The details of linear systems in arbitrary dimensions follow the patterns of potential eigenvalues – their multiplicities and complexities in full display. The halcyon days of trace and determinant fade, and the hard work of factoring and eigendecomposition proceeds.

This rigorous display of linear systems is more than a precursor: many of the serious and useful applications of this Volume are to systems with little-to-no nonlinearity. The nonlinear lies in the wake of the linear, thanks to the usual tools of Taylor. The removal of domain constraints to lines, circles, and planes gives rise to much more than equilibrium and periodic orbits. Stranger invariant sets haunt a few of the examples in this Volume, their forms as yet inchoate and awaiting their turn for attention.

MATRIX SOLUTIONS

Recall from the last Volume our observation that the solutions to linear systems with 2-by-2 matrices work just as well with k -by- k .

DISCRETE TIME: a discrete-time linear system on $\mathbb{R}^k$ is a dynamical system of the form $E\mathbf{x} = A\mathbf{x}$, where $\mathbf{x} \in \mathbb{R}^k$ is a vector variable. This has explicit solution $\mathbf{x}(n) = A^n \mathbf{x}_0$, where $\mathbf{x}_0 = \mathbf{x}(0)$ is the initial condition.

CONTINUOUS TIME: a continuous-time linear system on $\mathbb{R}^k$ is a dynamical system of the form $D\mathbf{x} = A\mathbf{x}$. Given an initial condition $\mathbf{x}_0 = \mathbf{x}(0)$, this first-order system of ODEs has solution $\mathbf{x}(t) = e^{At} \mathbf{x}_0$, where e^{At} is the **matrix exponential**,

$$e^{At} = \exp(At) = \sum_{n=0}^{\infty} \frac{(At)^n}{n!} = I + At + \frac{1}{2!} A^2 t^2 + \frac{1}{3!} A^3 t^3 + \dots$$

SIMPLE EIGENVALUES

The simplest possible case – increasingly rare with increasing dimension – is that of real eigenvalues $\{\lambda_i\}_1^k$ with independent eigenvectors $\{\mathbf{v}_i\}_1^k$, in which case the matrix is diagonalizable:

$$AV = V\Lambda \quad : \quad V = [\mathbf{v}_1 \mid \mathbf{v}_2 \mid \dots \mid \mathbf{v}_k] \quad : \quad \Lambda = \begin{bmatrix} \lambda_1 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \lambda_k \end{bmatrix}.$$

Powers and exponentials of A are computed via conjugation with V as

$$A^n = V\Lambda^n V^{-1} \quad : \quad e^{At} = V e^{\Lambda t} V^{-1}.$$

Powers and exponentials of a diagonal matrix being elementary, these systems are completely solved in this simplistic case. The issues of repeated and complex (& repeated complex) roots will recur; the next Chapter deals with this using the *Jordan Canonical Form*. A preview of that fuller story forms the remainder of this Chapter.

BASIS SOLUTIONS

Not all linear systems present themselves in matrix form. The generalization of the 2nd order systems from Volume 2 (so useful in oscillators) is a k^{th} order linear autonomous system. Such is expressible in operator notation as

$$\left(\sum_{i=0}^k a_i D^i \right) x = 0 \quad : \quad \left(\sum_{i=0}^k a_i E^i \right) x = 0 ,$$

in continuous and discrete time respectively. The **characteristic polynomial** of both polynomial evolution operators is given by

$$\sum_{i=0}^k a_i \lambda^i = 0 .$$

The roots of this are the eigenvalues of the matrix obtained by converting the k^{th} order system to a 1st order system on a vector.

EXAMPLE: The following fourth-order ODE has a conveniently factored polynomial differential operator:

$$\frac{d^4 x}{dt^4} + 3 \frac{d^3 x}{dt^3} - 4 \frac{dx}{dt} = 0 \quad \Rightarrow \quad D(D-1)(D+2)^2 x = 0 .$$

The characteristic polynomial, $\lambda^4 + 2\lambda^3 - 4\lambda^2 - 2\lambda$ has roots $\lambda = 0, 1, -2, -2$. One should check that these roots match the eigenvalues (with multiplicities) of the equivalent 1st-order matrix system:

$$D\mathbf{u} = A\mathbf{u} \quad : \quad \mathbf{u} = \begin{pmatrix} x \\ Dx \\ D^2x \\ D^3x \end{pmatrix} \quad : \quad A = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & -4 & 0 & 3 \end{bmatrix} .$$

For any linear n^{th} -order system on a variable x , the general solution is given as a linear combination of n **basis solutions**, denoted ϕ_i . There are several cases, depending on the nature and multiplicities of the roots (that is, the eigenvalues) and the nature of time. Repeated roots contribute additional basis functions, so that there is a spanning set of k solutions in total.

CONTINUOUS TIME:

- ▷ real λ : $\phi = e^{\lambda t}$
- ▷ real λ , multiplicity j : $\phi_1 = e^{\lambda t}$, $\phi_2 = te^{\lambda t}$, ..., $\phi_j = t^{j-1}e^{\lambda t}$
- ▷ complex $\alpha \pm i\beta$: $\phi_1 = e^{\alpha t} \cos \beta t$, $\phi_2 = e^{\alpha t} \sin \beta t$
- ▷ complex $\alpha \pm i\beta$, multiplicity j : $\phi_1 = e^{\alpha t} \cos \beta t$, ..., $\phi_{2j} = t^{j-1}e^{\alpha t} \sin \beta t$

In the case of discrete time with complex roots, then, as in Volume 2, we will want to convert to polar form $\alpha + i\beta = re^{i\theta}$.

DISCRETE TIME:

- ▷ real λ : $\phi = \lambda^n$
- ▷ real λ , multiplicity j : $\phi_1 = \lambda^n$, $\phi_2 = n\lambda^n$, ..., $\phi_j = n^{j-1}\lambda^n$
- ▷ complex $re^{\pm i\theta}$: $\phi_1 = r^n \cos n\theta$, $\phi_2 = r^n \sin n\theta$
- ▷ complex $re^{\pm i\theta}$, multiplicity j : $\phi_1 = r^n \cos n\theta$, ..., $\phi_{2j} = n^{j-1}r^n \sin n\theta$

In any case above, the *multiplicity* $j > 1$ is used to encode the number of repeated eigenvalues without independent eigenvectors. The general solution to a k^{th} order autonomous system is a linear combination of basis solutions,

$$x = \sum_{i=1}^k C_i \phi_i ,$$

where the constants C_i will require k bits of data (such as an initial condition and first $k - 1$ derivatives of $x(t)$ or the first k terms of x_n) to be uniquely determined.

EXAMPLE: In the previous example of a fourth-order ODE, the eigenvalues are $\lambda_1 = 0$, $\lambda_2 = 1$, and $\lambda_{3,4} = -2$; thus, the general solution is a linear combination of basis functions:

$$x(t) = C_1 + C_2 e^t + C_3 e^{-2t} + C_4 t e^{-2t} .$$

Determination of the unknown coefficients C_i depends on knowledge of four initial or boundary conditions.

EXAMPLE: Consider the following third-order linear recurrence relation – a generalization of the Fibonacci recurrence of Chapter 13:

$$G_n = G_{n-1} + G_{n-3} \quad : \quad G_0 = G_1 = 0 ; G_2 = 1 .$$

What will happen? Will the numbers grow monotonically or oscillate, and at what rates? The characteristic polynomial is $\lambda^3 - \lambda^2 - 1 = 0$, which has roots equal to $\lambda_1 \approx 1.466$ and $\lambda_{2,3} \approx -0.233 \pm 0.793 i \approx 0.826 e^{\pm 1.86 i}$. Using basis solutions, the general expression for G_n is given by:

$$G_n \approx C_1 (1.466)^n + (0.826)^n (C_2 \cos(1.86 n) + C_3 \sin(1.86 n)) .$$

The three constants are obtained from the initial conditions (with some effort). One notes that these not-so-Fibonacci numbers grow at a rate which approaches the not-so-mysterious ratio $G_n/G_{n-1} \rightarrow \lambda_1$ as $n \rightarrow \infty$.

DOMINANCE

Recall from Chapter 12 the notion of a **dominant eigenvalue**, which, in 2-D, measures the larger of the two (real distinct) eigenvalues (in real or absolute value, depending on continuous or discrete time). In n -D the definitions are somewhat more refined:

- ▷ *Continuous time*: λ_1 is dominant if $\Re(\lambda_1) > \Re(\lambda_k)$ for all $k > 1$.
- ▷ *Discrete time*: λ_1 is dominant if $|\lambda_1| > |\lambda_k|$ for all $k > 1$.

Here, modulus $|\cdot|$ is used in discrete time (modulus being the complex generalization of the absolute value). Only a real, non-repeated eigenvalue can be dominant (as complex eigenvalues come in conjugate pairs): it makes for a good exercise to modify the definition to accommodate complex pairs.

If λ_1 is the dominant eigenvalue, then, in the case of all real, distinct eigenvalues, the solutions to a linear system have the following asymptotics:

$$\begin{aligned} \mathbf{x}(n) &= \sum_{j=1}^K C_j \lambda_j^n \mathbf{v}_j = \lambda_1^n \left(C_1 \mathbf{v}_1 + C_2 \left(\frac{\lambda_2}{\lambda_1} \right)^n \mathbf{v}_2 \right) \longrightarrow C_1 \lambda_1^n \mathbf{v}_1 \\ \mathbf{x}(t) &= \sum_{j=1}^K C_j e^{\lambda_j t} \mathbf{v}_j = e^{\lambda_1 t} \left(C_1 \mathbf{v}_1 + \sum_{j=2}^K C_j e^{(\lambda_j - \lambda_1)t} \mathbf{v}_j \right) \longrightarrow C_1 e^{\lambda_1 t} \mathbf{v}_1 \end{aligned}$$

The situation is slightly more complicated with repeated or complex eigenvalues, but the long-term asymptotics follow the dominant eigenvalue term. It is a worthwhile exercise to show that in all such cases, as in the 2-D case:

- ▷ The asymptotic growth rate is determined by the dominant eigenvalue.

- ▷ The asymptotic ratio of terms is determined by the ratios of components of the dominant eigenvector.

(This is best proved using the Jordan Canonical Form of the next chapter.)

INPUT-OUTPUT MODELS

There is a broad class of linear/affine models in economics called *input-output* (or *Leontief models*, after their Nobel-prize-winning creator). Imagine an economic system partitioned into K sectors (e.g., robotics, semiconductors, AI, etc.), each reliant on the others for economic activity. Each sector produces a fixed type of output, which is then fed to other sectors for them to generate their respective outputs. Assume that economic production periods are discrete, based on time n . In the simplest setting, one has a discrete-time model that is purely linear: the vector $\mathbf{x} = (x_i)$ of economic outputs satisfies $E\mathbf{x} = A\mathbf{x}$ for some square matrix A , whose eigenvalues and eigenvectors are important for long-term behavior of the economy.

EXAMPLE: Consider an economy divided into five sectors corresponding to agriculture, manufacturing, construction, transportation, and services. Let $\mathbf{x}(n)$ denote the vector of five sector outputs at year n . A simplistic input-output model relating the outputs of these sectors is given by the following linear system:

$$E\mathbf{x} = A\mathbf{x} : A = \frac{1}{10} \begin{bmatrix} 4 & 1 & 1 & 2 & 1 \\ 3 & 5 & 2 & 1 & 1 \\ 1 & 1 & 4 & 1 & 2 \\ 2 & 2 & 1 & 3 & 1 \\ 1 & 4 & 1 & 1 & 5 \end{bmatrix}.$$

This A has eigenvalues $\lambda_1 \approx 1.035$, $\lambda_2 \approx 0.334$, $\lambda_{3,4} \approx 0.276 \pm .138i$, and $\lambda_5 \approx 0.178$. The corresponding eigenvectors are given by the following:

$$\mathbf{v}_1 \approx \begin{pmatrix} 0.59 \\ 0.88 \\ 0.64 \\ 0.62 \\ 1.00 \end{pmatrix} : \mathbf{v}_2 \approx \begin{pmatrix} -3.2 \\ -0.54 \\ 5.79 \\ -2.1 \\ 1.00 \end{pmatrix} : \mathbf{v}_{3,4} \approx \begin{pmatrix} -0.02 \pm 0.15i \\ -0.48 \pm 0.53i \\ -0.30 + 0.82i \\ 0.004 \pm 0.22i \\ 1.00 \end{pmatrix} : \mathbf{v}_5 \approx \begin{pmatrix} 1.16 \\ -0.65 \\ -0.61 \\ -1.15 \\ 1.00 \end{pmatrix}.$$

There is thus a dominant eigenvalue/eigenvector pair $(\lambda_1, \mathbf{v}_1)$, and over time, this economic system will grow at a rate of $\lambda_1 \approx 1.035 = 3.5\%$ per year. Services (x_5) will become the largest sector, comprising $1.00/|\mathbf{v}_1| \approx 58.5\%$ of the economy. The complex eigenvalue pair λ_3, λ_4 implies some oscillatory behavior. These oscillations die out over time, since these eigenvalues have modulus less than one. Whatever oscillatory behavior exists impacts agriculture and transportation the least, based on the relative sizes of the components of these eigenvectors.

CHAPTER 25 : THE JORDAN FORM

MATRIX FACTORIZATION is a core theme in Linear Algebra. The intricacies of solutions to linear systems are patterned after the factorization which decouples a system as much as possible. This optimal unchaining is captured by the normal form of Jordan.

DIAGONALIZATION REDUX

As in the 2-D case, the principal technique to solve a larger linear system is to decouple the system into subsystems of as small a dimension as possible. In the optimal case, the system is decomposable into 1-dimensional pieces. This corresponds to a factorization via an eigenbasis as per the previous chapter:

$$A = \Lambda \Lambda^{-1} \quad : \quad V = [\mathbf{v}_1; \mathbf{v}_2; \dots; \mathbf{v}_k] \quad : \quad \Lambda = \begin{bmatrix} \lambda_1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \lambda_k \end{bmatrix}.$$

Alas, such is not always possible. In 2-D, the simplest *forms* of matrices with repeated (λ) or complex ($\alpha \pm i\beta$) eigenvalues were, respectively,

$$R = \lambda I + N = \begin{bmatrix} \lambda & 1 \\ 0 & \lambda \end{bmatrix} \quad : \quad C = \alpha I + \beta J = \begin{bmatrix} \alpha & -\beta \\ \beta & \alpha \end{bmatrix},$$

where the special 2-by-2 matrices I, J , & N (see Chapters 11 and 12) can simplify the notation somewhat.

In an n -dimensional setting, the best that one can hope for in general is a *block diagonal* matrix, where simplified subsystems are mutually decoupled. It is easier to draw a cartoon of a block diagonal matrix than it is to describe it accurately, but a constructive approach is not too opaque: begin with a diagonal matrix (all off-diagonal terms are zero) and “blow up” each entry to a matrix, making sure that (1) the sizes are compatible, and (2) the off-diagonal terms are blown-up to zero-matrices. The block diagonal matrices arising as Jordan forms are both square and all the *Jordan blocks* along the diagonal are also square (though of potentially different sizes).

The *Jordan Classification Theorem* asserts that for any square matrix A there is a factorization of the form $A = VJV^{-1}$, where J is a Jordan block diagonal matrix and V is a change of basis matrix. The Jordan form J is unique up to the ordering of the blocks J_k along the diagonal: that is, unique up to permuting coordinates.

COMPLEX JORDAN BLOCKS

The question remains: what are the Jordan blocks? A typical linear algebra course would present factorization over the complex field $\mathbb{C}$, in which case the Jordan blocks are simple to describe. Each is a square matrix of the form

$$J_* = \lambda I + N = \begin{bmatrix} \lambda & 1 & 0 & \cdots & 0 & 0 \\ 0 & \lambda & 1 & 0 & \cdots & 0 \\ 0 & 0 & \lambda & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & \ddots & 1 & 0 \\ \vdots & \vdots & \ddots & \ddots & \lambda & 1 \\ 0 & 0 & 0 & \cdots & 0 & \lambda \end{bmatrix}.$$

Here, I is an identity matrix and N is the correspondingly-sized zero matrix with ones on the superdiagonal. The simplest 1-by-1 block is an eigenvalue with an independent eigenvector (potentially complex). Repeated eigenvalues sharing a single eigenvector induce a block J_* with ones on the superdiagonal. Note that (1) this permits complex values for λ , at the expense of allowing for complex-valued eigenvectors; (2) this description subsumes the 2-D case of Volume 2.

For purposes of linear algebra, this description is complete. However, in the context of solving dynamical systems with real variables, one is left with an unsatisfying complex-valued solution to a real-valued problem. The user has two choices: to derive a complex-valued solution, then split into real and imaginary parts; or propagate the splitting back to the Jordan form. It is the latter that was done in Volume 2, for concreteness. Extending this to arbitrary dimensions is worthwhile, but a bit more work.

REAL JORDAN BLOCKS

Over the reals, Jordan blocks come in four types:

- ▷ The 1-by-1 block $[\lambda]$ for a *real, independent* eigenvalue;
- ▷ The 2-by-2 block associated with a *complex* pair $\alpha \pm \beta i$

$$C = \begin{bmatrix} \alpha & -\beta \\ \beta & \alpha \end{bmatrix};$$

- ▷ The j -by- j block $[\lambda]$ for a *real, repeated* non-independent eigenvalue

$$R = \begin{bmatrix} \lambda & 1 & 0 & \cdots & 0 & 0 \\ 0 & \lambda & 1 & 0 & \cdots & 0 \\ 0 & 0 & \lambda & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & \ddots & 1 & 0 \\ \vdots & \vdots & \ddots & \ddots & \lambda & 1 \\ 0 & 0 & 0 & \cdots & 0 & \lambda \end{bmatrix};$$

These are all best written in condensed form as $C = \alpha I + \beta J$ and $R = \lambda I + N$. A typical Jordan form will have several different types of such elementary blocks arranged along the (block) diagonal.

In answer to the reader's queries: (1) yes, one must deal with the case of repeated complex eigenvalues; (2) no, it is not pleasant. The author is unaware of an application in which this case makes a significant contribution (and does not necessarily wish for an end to his ignorance). Nevertheless, it must be addressed. The simplest approach is to use a compressed matrix-of-2-by-2-blocks:

- ▷ The $2j$ -by- $2j$ block $[\lambda]$ for a *complex* pair $\alpha \pm \beta i$, *repeated* j times, sharing a common eigenspace

$$\begin{bmatrix} C & I & 0 & \cdots & 0 & 0 \\ 0 & C & I & 0 & \cdots & 0 \\ 0 & 0 & C & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & \ddots & I & 0 \\ \vdots & \vdots & \ddots & \ddots & C & I \\ 0 & 0 & 0 & \cdots & 0 & C \end{bmatrix}.$$

Each entry of the above matrix is a 2-by-2 block, with C depending on α and β as before. The reader will hopefully not encounter this block type in the wild. It is helpful to have real blocks for real solutions, but the reader would not be wrong to sympathize with the mathematicians' complex approach.

EXAMPLE: What can be said if a matrix A has Jordan Canonical Form equal to

$$V^{-1}AV = J = \begin{bmatrix} 2 & 1 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 & 0 \\ 0 & 0 & 2 & 0 & 0 \\ 0 & 0 & 0 & 2 & 1 \\ 0 & 0 & 0 & -1 & 2 \end{bmatrix}?$$

This 5-by-5 matrix A has a real eigenvalue 2 with multiplicity three, but it is a bit more complicated, since one eigenvalue has its own independent eigenspace, and the other two

copies of that eigenvalue share one eigenspace; hence the 2-by-2 block in the upper left corner. Although there are 2's all along the diagonal, the remaining pair of eigenvalues is the complex conjugate pair $2 \pm i$. In a linear dynamical system there would be a great deal of instability, with some oscillatory behavior as well. (Note that in general, one should specify continuous versus discrete time: in this instance, they have similar dynamics!)

POWERS & EXPONENTS

Assuming a set of (generalized) eigenvectors $\{\mathbf{v}_i\}_1^k$ give column vectors of a matrix V by which one transforms a matrix A into its Jordan canonical form, the problems of computing explicit solutions to $D\mathbf{x} = A\mathbf{x}$ or $E\mathbf{x} = A\mathbf{x}$ become simple. As in the 2-D case, powers and exponentials of A are computed via conjugation with V as $A^n = V\Lambda^n V^{-1}$ and $e^{At} = V e^{\Lambda t} V^{-1}$. Block-diagonal matrices retain their block structure under powers (and thus exponentials), and so, it remains only to compute powers and exponentials of Jordan blocks. The simplest case of 1-by-1 blocks is trivial; the case on complex 2-by-2 blocks was derived in Chapter 12. The case of a real, repeated eigenvalue of multiplicity j is more interesting: recalling the matrix $R = \lambda I + N$ above, a derivation similar to that of the 2-by-2 case in Chapter 12 reveals the following:

$$R^n = \sum_{i=0}^n \binom{n}{i} \lambda^{n-i} N^i = \begin{bmatrix} \lambda^n & n\lambda^{n-1} & \binom{n}{2}\lambda^{n-2} & \dots & \binom{n}{j-1}\lambda^{n-j+1} \\ 0 & \lambda^n & n\lambda^{n-1} & \ddots & \vdots \\ 0 & 0 & \lambda^n & \ddots & \binom{n}{2}\lambda^{n-2} \\ \vdots & \vdots & \vdots & \ddots & n\lambda^{n-1} \\ 0 & 0 & 0 & \dots & \lambda^n \end{bmatrix},$$

$$e^{Rt} = e^{(\lambda t)I} \sum_{i=0}^{\infty} \frac{1}{i!} N^i t^i = \begin{bmatrix} e^{\lambda t} & te^{\lambda t} & \frac{1}{2!}t^2e^{\lambda t} & \dots & \frac{1}{(j-1)!}t^{j-1}e^{\lambda t} \\ 0 & e^{\lambda t} & te^{\lambda t} & \ddots & \vdots \\ 0 & 0 & e^{\lambda t} & \ddots & \frac{1}{2!}t^2e^{\lambda t} \\ \vdots & \vdots & \vdots & \ddots & te^{\lambda t} \\ 0 & 0 & 0 & \dots & e^{\lambda t} \end{bmatrix}.$$

Once again, the Binomial Theorem (in discrete time) and the definition of the exponential function (in continuous time) gives the true reason for the strange coefficients in the matrix entries. One should compare these powers and exponentials with the basis-solution approach and reconcile the similarities and differences. The case of complex repeated roots is too much to appear in print.

To reiterate: once the Jordan Canonical Form is known (to obtain this, as well as its change of coordinates is no mean feat!) one can present a general solution to any linear problem by exploiting the block structure:

$$J^n = \begin{bmatrix} J_1^n & 0 & 0 & \dots & 0 \\ 0 & J_2^n & 0 & \ddots & \vdots \\ 0 & 0 & J_3^n & \ddots & 0 \\ \vdots & \vdots & \vdots & \ddots & 0 \\ 0 & 0 & 0 & \dots & J_k^n \end{bmatrix} : e^{Jt} = \begin{bmatrix} e^{J_1 t} & 0 & 0 & \dots & 0 \\ 0 & e^{J_2 t} & 0 & \ddots & \vdots \\ 0 & 0 & e^{J_3 t} & \ddots & 0 \\ \vdots & \vdots & \vdots & \ddots & 0 \\ 0 & 0 & 0 & \dots & e^{J_k t} \end{bmatrix}.$$

CHAPTER 26 : LINEARIZATION & STABILITY

LINEAR ALGEBRA fuels the analysis of nonlinear systems, thanks to the technique of linearization about equilibria. There are very few local surprises in the passage from 2-D to arbitrary dimensions, apart from the increased complexities encoded in the Jordan form. This increase in dimensional complexity does afford more realistic models of physical and social systems.

LINEARIZATION

Linearization about an equilibrium proceeds exactly according to the pattern set in the story of Volume 2. Given $D\mathbf{x} = F(\mathbf{x})$ or $E\mathbf{x} = F(\mathbf{x})$ where $\mathbf{x} = (x_i)$ is a vector depending on time (continuous t or discrete n), one begins by searching for the equilibria, where $D\mathbf{x} = F(\mathbf{x}) = \mathbf{0}$ or $E\mathbf{x} = F(\mathbf{x}) = \mathbf{x}$ respectively. This can be very difficult in practice, since the algebra is nonlinear. Once an equilibrium is found, local analysis proceeds with linearizing the right hand side F .

The derivative of F at input $\mathbf{x} = \mathbf{a}$, $[DF]_{\mathbf{a}}$, is a linear transformation whose entries are the partial derivatives of the outputs with respect to the inputs. This is most often represented as a square matrix. Setting $\mathbf{h} = \mathbf{x} - \mathbf{a}$ as a perturbation to an equilibrium, the local dynamics on this perturbation are, as in Chapter 15,

$$\begin{aligned} D\mathbf{h} &= D\mathbf{x} - D\mathbf{a} = F(\mathbf{a} + \mathbf{h}) - F(\mathbf{a}) = [DF]_{\mathbf{a}}\mathbf{h} + O(|\mathbf{h}|^2) \approx [DF]_{\mathbf{a}}\mathbf{h} \\ E\mathbf{h} &= E\mathbf{x} - E\mathbf{a} = F(\mathbf{a} + \mathbf{h}) - \mathbf{a} = [DF]_{\mathbf{a}}\mathbf{h} + O(|\mathbf{h}|^2) \approx [DF]_{\mathbf{a}}\mathbf{h} \end{aligned}$$

These linearized systems are explicitly solvable:

$$\mathbf{h}(t) \approx \exp([DF]_{\mathbf{a}}t) \mathbf{h}_0 \quad : \quad \mathbf{h}_n \approx ([DF]_{\mathbf{a}})^n \mathbf{h}_0.$$

One hopes that the linearized dynamics are a good approximation to the full nonlinear dynamics: the next section explores this.

EXAMPLE: Consider the following differential equations.

$$\frac{dx}{dt} = 2y - 5x \quad : \quad \frac{dy}{dt} = 7x - y - xz \quad : \quad \frac{dz}{dt} = xy - 3z$$

This has an equilibrium at the origin, since the right hand side vanishes there. Computing the derivative and evaluating at the origin yields:

$$[DF]_{\mathbf{0}} = \begin{bmatrix} -5 & 2 & 0 \\ 7 & -1 & -x \\ y & x & -3 \end{bmatrix}_{\mathbf{0}} = \begin{bmatrix} -5 & 2 & 0 \\ 7 & -1 & 0 \\ 0 & 0 & -3 \end{bmatrix}$$

Note the built-in block-diagonal structure. That is good fortune, since it allows us to avoid explicitly computing eigenvalues. In the linearized dynamics, the z -axis is invariant and stable; the (x, y) plane is likewise invariant but not stable. Since the determinant of the 2-by-2 block is negative, the dynamics in the (x, y) plane exhibit a saddle.

STABILITY

The recurring lesson of this text is that eigenvalues are primal. Though the example above was serendipitous, it is no longer possible in general to use a quickhack approach such as the trace/determinant of the derivative to predict what type of equilibrium one has: eigenvalues, rather than equilibria, are what one classifies. The following language will be used throughout the remainder of this text.

- ▷ DISCRETE TIME : An eigenvalue λ is said to be
 - STABLE if $|\lambda| < 1$ (modulus in the case of a complex λ)
 - NEUTRAL if $|\lambda| = 1$ (modulus in the case of a complex λ)
 - UNSTABLE if $|\lambda| > 1$ (modulus in the case of a complex λ)

- ▷ CONTINUOUS TIME : An eigenvalue λ is said to be
 - STABLE if $\Re(\lambda) < 0$ (real part in the case of a complex λ)
 - NEUTRAL if $\Re(\lambda) = 0$ (real part in the case of a complex λ)
 - UNSTABLE if $\Re(\lambda) > 0$ (real part in the case of a complex λ)

EXAMPLE: In the example above, the equilibrium at the origin has two stable eigenvalues and one unstable eigenvalue. This equilibrium could be called a *saddle*, though it is more informative to give the dimensions of the stable (1) and unstable (2) eigenspaces.

In general, the stable/unstable dichotomy of Volume 1 is less useful in higher dimensional systems. Our convention (not universal!) is as follows. An equilibrium will be called **stable** only if *all* eigenvalues are stable; if *any* eigenvalue is unstable, the equilibrium will be called **unstable**. The presence of neutral eigenvalues *might* destabilize an equilibrium all of whose other eigenvalues are stable; a single unstable eigenvalue leads to an unstable equilibrium, independent of whether other eigenvalues are stable or neutral.

HYPERBOLICITY & CONJUGACY

Recall that an equilibrium is **hyperbolic** if there are no neutral eigenvalues. The differences between 2-D and higher dimensional systems are minimal when it comes to trusting linearization. The principal result used to guarantee fidelity is and remains the following.

HARTMAN-GROBMAN THEOREM: On a sufficiently small neighborhood of a hyperbolic fixed point $\mathbf{a}$, the nonlinear dynamics of F (assumed C^1) are topologically conjugate to the dynamics of the linearized system $[DF]_{\mathbf{a}}$.

There are a few points worth reviewing.

- ▷ It is not enough to say “centers are suspect” since we cannot attach names to all possible types of equilibria in n -D. One must work with stable, unstable, and neutral eigenvalues.
- ▷ In n -D, hyperbolic equilibria are increasingly rare; the more eigenvalues, the more chances that one or more are neutral.
- ▷ This result is *local*; linearization provides local information only.
- ▷ There are subtleties bound with the notion of topological conjugacy.

As with the 2-D case, the full Hartmann-Grobman Theorem gives much more information than mere local fidelity. Eigenvector directions are also preserved by the topological conjugacy, providing more structured details. See Chapter 31 for more information.

The remainder of this chapter will consist of examples of our story in action – finding and classifying equilibria.

THE HEART OF STABILITY

The *Zeeman model* of electrochemical activity in the heart is a continuous-time model in 3-D. The relevant variables are the sodium conductance $x(t)$, the potassium conductance $y(t)$, and a control potential across the cell membrane $z(t)$. The system is given by the equations:

$$\frac{d}{dt} \begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} -\frac{1}{\epsilon}(x^3 + xy + z) \\ -2x - 2y \\ -y - 1 \end{pmatrix} : \quad 0 < \epsilon \ll 1.$$

Here, ϵ is a small positive parameter. Setting the right hand side equal to zero reveals a single equilibrium at $(1, -1, 0)$. Evaluating the derivative here yields

$$\begin{bmatrix} -\frac{1}{\epsilon}(3x^2 + y) & -\frac{x}{\epsilon} & -\frac{1}{\epsilon} \\ -2 & -2 & 0 \\ 0 & -1 & 0 \end{bmatrix}_{1,-1,0} = \begin{bmatrix} -\frac{2}{\epsilon} & -\frac{1}{\epsilon} & -\frac{1}{\epsilon} \\ -2 & -2 & 0 \\ 0 & -1 & 0 \end{bmatrix}.$$

Solving for the eigenvalues is by no means straightforward – even with computer algebra, the answer is unhelpfully precise. It is best to think asymptotically with respect to the small parameter ϵ . To leading order, the eigenvalues are

$$\lambda_1 \approx -\frac{2}{\epsilon} : \lambda_{2,3} \approx \frac{1}{2}(-1 \pm i\sqrt{3}).$$

The equilibrium is therefore clearly stable, both linearly and nonlinearly, thanks to Hartman-Grobman. The dominant eigenvalue λ_1 (remember, $\epsilon \ll 1$) pulls all initial conditions down to a 2-D surface on which there is a spiral sink.

This is an interesting example from multiple perspectives. One concludes that (nearby) solutions have a damped oscillatory behavior. The role that the small parameter ϵ plays is particularly significant. The dynamics on the x variable are much faster than those of the other two: this is an example of a *fast-slow* system. As well, the model is more general than that of electrical activity: the general cubic governing the dynamics of $x(t)$ is applicable to a large number of systems near a cusp singularity (recall Chapter 9 and note the presence of Zeeman both there and here).

THE INTERMEDIATE AXIS THEOREM

Consider a solid body in 3-D which is spinning about an axis through its center of mass. Ignoring gravity and air-friction (!) one has a freely-spinning object which would seem to be tracing out a periodic orbit in some continuous-time dynamical system. Set (x_1, x_2, x_3) a coordinate frame whose origin is at the center of mass and whose axes are aligned with the three *principal axes* of the massive body (these are eigenvectors of the body's *inertia matrix*). The three principal moments of inertia I_1, I_2, I_3 (the eigenvalues of said matrix) represent the resistance to rotation about the three principal axes. Assume that the body has its three principal moments all distinct, arranged so that $I_1 > I_2 > I_3$. A good example is a rectangular box with three dimensions distinct, like a typical book.

One can observe by throwing such an object up in the air with a spin that rotation about the first and third principal axes feels somewhat stable, in the sense that a good spin about that axis seems to persist while in the air. Rotation about the middle (or *intermediate*) axis or any non-axis-aligned rotation leads to apparently floppy or erratic-looking rotation while in the air. [Please do not throw good books in the air.] What is happening and why?

The so-called *Euler equations* govern the three components (ℓ_1, ℓ_2, ℓ_3) of the angular momentum vector $\mathbf{L}$ as follows:

$$\frac{d\mathbf{L}}{dt} = \frac{d}{dt} \begin{pmatrix} \ell_1 \\ \ell_2 \\ \ell_3 \end{pmatrix} = \begin{pmatrix} \frac{I_2 - I_3}{I_2 I_3} \ell_2 \ell_3 \\ \frac{I_3 - I_1}{I_1 I_3} \ell_1 \ell_3 \\ \frac{I_1 - I_2}{I_1 I_2} \ell_1 \ell_2 \end{pmatrix}.$$

The usual approach seems to fail here, as every point on the three axes is an equilibrium, necessarily degenerate. What is to be done? The absence of friction in this model is a clue that the dynamics are conservative. Conservation of angular momentum means that $\Phi = \ell_1^2 + \ell_2^2 + \ell_3^2$ is constant. This means that we can restrict attention to a radius C sphere of constant angular momentum. On this sphere, there are six equilibria, where the axes pierce the sphere, at $(\pm C, 0, 0)$, $(0, \pm C, 0)$, and $(0, 0, \pm C)$. The linearized dynamical system at the three pairs of equilibria are, respectively,

$$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & \pm \frac{I_3 - I_1}{I_1 I_3} C \\ 0 & \pm \frac{I_1 - I_2}{I_1 I_2} C & 0 \end{bmatrix} : \begin{bmatrix} 0 & 0 & \pm \frac{I_2 - I_3}{I_2 I_3} C \\ 0 & 0 & 0 \\ \pm \frac{I_1 - I_2}{I_1 I_2} C & 0 & 0 \end{bmatrix} : \begin{bmatrix} 0 & \pm \frac{I_2 - I_3}{I_2 I_3} C & 0 \\ \pm \frac{I_3 - I_1}{I_1 I_3} C & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

These are all (up to a permutation of coordinates) block-diagonal, with a 1-by-1 block with eigenvalue zero (corresponding to the degenerate eigenvalue) and a 2-by-2 block with off-diagonal terms depending on the three principal moments $I_1 > I_2 > I_3$. Since the 2-by-2 blocks determine the dynamics on the tangent plane to the sphere at the equilibrium, we can use the eigenvalues (or determinants!) to classify the equilibria on the sphere. For the first and third pairs above – corresponding to rotation about the major and minor axes – the eigenvalues are *pure imaginary* and the rotation is semi-stable (not unstable, at least). However, for rotation about the intermediate axis – given by the middle matrix above – the equilibrium is a saddle, since $I_1 > I_2 > I_3$. This is the reason why such rotations appear to flip about. Conservation means that the predicted centers are true centers and that all orbits (except those between the two saddles) are periodic. The full picture of how the local classification fits into global behavior gives a fascinating example of the power of eigenvalues to predict behavior.

DYNAMO DYNAMICS

Three-dimensional continuous-time models can push us beyond what we can currently handle. Consider the following classical 1958 model by Rikitake for the geomagnetic dynamics of the Earth's core. Beginning with a simplistic reduction to a pair of magnetic discs with currents x and y spinning symmetrically with angular momentum z , the follow set of differential equations emerges:

$$\frac{d}{dt} \begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} -bx + yz \\ -ax - by + xz \\ 1 - xy \end{pmatrix}.$$

Here, a and b are positive parameters representing separation between the discs and resistive dissipation respectively. To find the equilibria, use the first two equations to eliminate z ; use the last equation to obtain $xy = 1$; then substitute this in to obtain $x^4 - ax^2/b = 1$. An application of the quadratic formula yields a pair of equilibria at

$$x = \pm \sqrt{\frac{a + \sqrt{a^2 + 4b}}{2b}} : y = \pm \sqrt{\frac{2b}{a + \sqrt{a^2 + 4b}}}.$$

Linearization and some unpleasant algebra yields the three eigenvalues $\lambda_1 = -2b$ and $\lambda_{2,3} = \pm 2i \left(1 + \frac{a}{2b}\right)^{1/4}$. The equilibria have one stable eigenspace pulling solutions down to 2-D surfaces on which the equilibria are predicted to be centers. Alas, Hartmann-Grobman is again of limited use here: we need more

data. Simulations seem to indicate that the equilibria are very weakly unstable in their center-like 2-D eigenspaces. However, more disturbing behavior is revealed. Solutions quickly converge to what seems like a 2-D surface. Near the equilibria, solutions spiral out ever-so-slowly, until, eventually, there is a dramatic flip from orbiting about one equilibrium to orbiting about the other. This is of importance in modelling the mysterious geomagnetic reversals our planet is sometimes subject to. It is as well a mystery at the foundations of dynamics: if orbits are neither running off to infinity nor converging to periodic orbits or equilibria, what else can happen? Without the Poincaré-Bendixson Theorem, we have no answer – yet. Resolution will have to wait until Volume 4.

It is at this cusp of chaotic dynamics that we pull back and return to safe ground. Our models in this chapter have focused on continuous-time systems, as these are somewhat easier to work with in three dimensions. Discrete-time systems (even “boring” linear systems) will be the centerpiece of our next class of models.

CHAPTER 27 : POPULATION COHORT MODELS

POPULATION MODELS go well-beyond the 2-species variants of Volume 2. This, the first set of serious examples of high-dimensional dynamical systems in this Volume, will use a broad set of tools from both continuous and discrete-time systems, linear and nonlinear. The technique of dominant eigenvalues and eigenvectors will mature and demonstrate its predictive powers.

AGE COHORT DYNAMICS

Population dynamics are often naturally discretized into *age cohorts*: the age variable is partitioned into equal-sized intervals. This may occur by year, as in the case of livestock or, with its own special terminology, students (freshman, sophomore, junior, senior). Human population at the scale of nations is often partitioned into cohorts by decade – “a child of the ‘60s”.

Given such a partition into k age cohorts, one packs the age cohort sizes into a population vector $\mathbf{p} = \mathbf{p}(n)$ which depends on a discrete time n , where the dynamics involves moving the clock ahead one time unit (year, decade, ...). In the simplest model, each cohort *graduates* to the next, giving the linear system:

$$\mathbf{p}(n) = \begin{pmatrix} p_1(n) \\ p_2(n) \\ p_3(n) \\ \vdots \\ p_k(n) \end{pmatrix} : \quad \mathbf{p}(n+1) = \begin{bmatrix} 0 & 0 & 0 & \cdots & 0 \\ 1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ \vdots & 0 & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 & 0 \end{bmatrix} \mathbf{p}(n) = \begin{pmatrix} 0 \\ p_1(n) \\ p_2(n) \\ \vdots \\ p_{k-1}(n) \end{pmatrix}.$$

The matrix, conveniently represented in block form as the zero matrix with a sub-diagonal identity, pushes each age cohort to the next. As there are no “births” and everyone “graduates” in the end, this matrix is nilpotent: *momento mori*.

A more realistic model would incorporate births. Such can be represented by replacing the zeros in the first row of the matrix above with nonnegative constants representing average reproduction by cohort:

$$\mathbf{p}(n+1) = \begin{bmatrix} r_1 & r_2 & r_3 & \cdots & r_k \\ 1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ \vdots & 0 & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 & 0 \end{bmatrix} \mathbf{p}(n).$$

These coefficients may be zero (e.g., no reproduction if too young or too old). Most realistic models would have a unimodal sequence of r_i terms. To improve the model yet further, one could replace the lower-left identity matrix with a diagonal matrix of coefficients $0 < s_i \leq 1$ representing *survival rates* – what proportion of the cohort graduates or survives to the next age.

$$\mathbf{p}(n+1) = \begin{bmatrix} r_1 & r_2 & r_3 & \cdots & r_k \\ s_1 & 0 & 0 & \cdots & 0 \\ 0 & s_2 & 0 & \cdots & 0 \\ \vdots & 0 & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & s_{k-1} & 0 \end{bmatrix} \mathbf{p}(n).$$

This model, with reproduction and survival coefficients, is known as the **Leslie model**. It is, strictly speaking, modeling the populations of reproductive females in the species, with the assumption that this is roughly half the total population size. One can imagine further improvements to account for additional complexities.

As a dynamical system, this hardly seems worth attention. There is one undesirable equilibrium at the origin. One expects the population to grow or

shrink over time, depending on the coefficients, with perhaps some interesting movements between cohorts. The truth is that linear systems should be by no means be disparaged, as they are the key to understanding nonlinear systems.

EIGENVALUE ISSUES

It is very much worthwhile to dig into the details of these cohort models, if for no other reason than to practice the linear-analysis portion of our story. Consider first the intermediate model above with reproduction but full survival: $s_i = 1$ for all i . The characteristic polynomial of this matrix is

$$\lambda^k - \sum_{i=1}^k r_i \lambda^{k-i} = 0.$$

LEMMA: When all the $r_i > 0$, this system has a strictly positive dominant eigenvalue λ_* along with an eigenvector whose entries are also strictly positive.

▷ *Proof*: Divide the characteristic polynomial by the highest power of λ to obtain a polynomial (known as a **Lotka-Euler equation**):

$$1 = \sum_{i=1}^k \frac{r_i}{\lambda^i} = f(\lambda).$$

As a function of a real variable, the right hand side $f(\lambda)$ has 1st derivative negative, is unbounded as $\lambda \rightarrow 0^+$, and vanishes as $\lambda \rightarrow \infty$. Continuity of f yields a unique positive root λ_* . The remaining roots must be of the form $\lambda = ae^{\pm i\theta}$ and can, with sufficient work, be shown to have $|a| < \lambda_*$ (using the fact that the real part of such a λ equals $a \cos \theta$ and $\cos \theta < 1$ since λ is not a positive real).

The eigenvector associated with the dominant eigenvalue by choosing the first entry to 1 can be shown, by the Lotka-Euler equation above, to be equal to

$$\mathbf{v}_* = \begin{pmatrix} 1 \\ \lambda_*^{-1} \\ \lambda_*^{-2} \\ \vdots \\ \lambda_*^{1-k} \end{pmatrix},$$

and all these entries are positive. ◁

This is significant in that the population stabilizes over time to a power distribution that depends very sensitively on the dominant eigenvalue λ_* , growing or decaying depending on which side of unity the dominant eigenvalue lies.

The above lemma holds in the case where the survival rates are non-uniform, so long as they are positive. In this case, the net survival coefficients $\sigma_i = s_1 s_2 \cdots s_i$ play a role, in that the dominant eigenvector can be chosen to be

$$\mathbf{v}_* = \begin{pmatrix} 1 \\ \sigma_1 \lambda_*^{-1} \\ \sigma_2 \lambda_*^{-2} \\ \vdots \\ \sigma_{k-1} \lambda_*^{1-k} \end{pmatrix}.$$

CASTES & CONSTRAINTS

Cohorts in population models can take other forms besides age. Consider the following model of a caste system (from the 1979 text of Luenberger), based on that of the Natchez tribe. The population is partitioned into four castes, rank-ordered from highest to lowest. Reproduction is regulated to mix populations, with the following two rules in effect:

1. Anyone not in the lowest caste must marry someone from the lowest caste.
2. Children are assigned to a caste which is the higher of the mother's caste or one less than the father's caste.

This complex set of constraints prompts a dynamical population model. Assume that populations are grouped into discrete-time generations, indexed by time n . Furthermore, let x_A, x_B, x_C, x_D represent the population sizes of the four castes as functions of n . Specifically, let these variables represent the number of families, assumed equal to the number of mothers or, equivalently, fathers. (Such assumptions are of course not modernly enlightened.)

		MOTHER			
		A	B	C	D
FATHER	A	X	X	X	B
	B	X	X	X	C
	C	X	X	X	D
	D	A	B	C	D

It is helpful to build a table of permissible pairings with offspring caste assignments in order to generate a linear dynamical system for population sizes from one generation to the next (assume total in-generation pairings). The derived model, assuming a reproduction rate λ , is given by

$$E \begin{pmatrix} x_A \\ x_B \\ x_C \\ x_D \end{pmatrix} = \lambda \begin{bmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ -1 & -1 & 0 & 1 \end{bmatrix} \begin{pmatrix} x_A \\ x_B \\ x_C \\ x_D \end{pmatrix}.$$

Arguing through this is a nontrivial exercise. The first row is simple: the only source of A-children is through an A-mother. The symmetry between mothers and fathers helps argue the second row, noting that B children come only from B mothers or A fathers. The third row is similar. The last row is difficult to grasp and is best written out as

$$x_D(n+1) = \lambda x_C(n) + \lambda(x_D(n) - x_A(n) - x_B(n) - x_C(n)),$$

the former term coming from the third row; the latter from the last column.

The long-term dynamics of this system is unfortunate. The eigenvalues are all λ and the 4-by-4 matrix is $(\lambda I + N)$, where N is a nilpotent matrix. The full solution is

$$x(n) = \lambda^n \begin{bmatrix} 1 & 0 & 0 & 0 \\ n & 1 & 0 & 0 \\ \frac{1}{2}n(n-1) & n & 1 & 0 \\ -\frac{1}{2}n(n+1) & -n & 0 & 1 \end{bmatrix} x(0) = \lambda^n \begin{pmatrix} x_A(0) \\ nx_A(0) + x_B(0) \\ \frac{1}{2}n(n-1)x_A(0) + nx_B(0) + x_C(0) \\ -\frac{1}{2}n(n+1)x_A(0) + -nx_B(0) + x_D(0) \end{pmatrix}.$$

As might have been guessed, one runs out of people in the lowest caste.

POPULATION CONTROLS

The previous caste example is worth careful consideration. An explicit solution produces an overabundance of people in castes B and C with a quadratic collapse of population size in caste D.

One is tempted to step in and solve the problem. The exhaustion of the lowest caste, may, with sufficient intervention, be stemmed. Assume that the mating restrictions are kept in place but that reproduction rates are controlled (through means ignored here, as in the real world). If one updates the table of permitted pairings to include reproduction rates α_i and λ_i as shown, then the resulting linear system becomes

		MOTHER			
		A	B	C	D
FATHER	A	X	X	X	α_1
	B	X	X	X	α_2
	C	X	X	X	α_3
	D	λ_1	λ_2	λ_3	λ_4

$$\begin{pmatrix} x_A \\ x_B \\ x_C \\ x_D \end{pmatrix} = \begin{bmatrix} \lambda_1 & 0 & 0 & 0 \\ \alpha_1 & \lambda_2 & 0 & 0 \\ 0 & \alpha_2 & \lambda_3 & 0 \\ -\lambda_4 & -\lambda_4 & \alpha_3 - \lambda_4 & \lambda_4 \end{bmatrix} \begin{pmatrix} x_A \\ x_B \\ x_C \\ x_D \end{pmatrix}.$$

This is not as simple to solve. Fortunately, the matrix, L , above, is triangular and the eigenvalues are $\lambda_1, \lambda_2, \lambda_3$, and λ_4 , all positive. Assuming they are distinct (our impetus for working with this system), there is a dominant eigenvalue whose associated eigenvector determines the long-term population distribution. The goal is to have such a dominant eigenvector whose entries are all positive: no caste is bred out of existence.

It is at this point where one must choose which caste to promote and which to inhibit. Given that we depleted the last caste – whose members must accommodate all other castes – it seems prudent to boost their number. If we set λ_4 to be maximal, then what eigenvector pairs with it? Unfortunately, the solution to $(L - \lambda_k I)\mathbf{v} = \mathbf{0}$ must have the first entry, $v_A = 0$, unless $k = 1$. Despite the observation that the A population caused a depletion of the D population and an excess in B and C , the only possible choice for a dominant eigenvalue is λ_1 – the A -caste must be encouraged to reproduce more.

This is not sufficient, as there are a number of subtle conditions on the reproduction rates that must be satisfied in order to have a dominant eigenvector with all entries positive: the interested reader will want to work through this as an exercise.

The fact that intuition is a poor guide in a simple, linear, deterministic model should give the reader pause in the face of temptation to pontificate. Whether demographic, economic, or ecological systems are under consideration, the long-term results of acquiescing to “*something must be done now*” arguments are often an unintended collapse.

INFECTION

Population cohorts based on age or generation are by no means exhaustive of applications. In epidemiology, one often works with cohorts in a fixed population size. Consider the popular SEIR model. In a well-mixed population of size N , there are four cohorts:

- ▷ SUSCEPTIBLE : of size S ;
- ▷ EXPOSED : of size E ;
- ▷ INFECTED : of size I ;
- ▷ RECOVERED : of size R .

There are numerous variations on this partition, with some models incorporating the possibility of reinfection, immunity, or death. The dynamics on these cohort sizes may be in continuous or discrete time.

EXAMPLE: The classical SIR model in continuous time is given by:

$$\begin{aligned}\frac{dS}{dt} &= -\frac{\beta}{N}IS \\ \frac{dI}{dt} &= \frac{\beta}{N}IS - \gamma I \\ \frac{dR}{dt} &= \gamma I.\end{aligned}$$

Here, $\beta, \gamma > 0$ are parameters, and N is a constant representing total population size. One should note that, although this is a nonlinear model, there are some similarities with the linear population cohort models of this chapter, including a notion of “graduation” to the next cohort. In the case of passage from susceptible to infected, this is proportional to the mixing between susceptible and infected agents – hence the quadratic term. In the case of passage from infected to recovered, there is no mixing involved: it is a linear function.

The dynamics of SIR and related cohort models in epidemiology is unexciting. The most interesting features from a dynamical perspective involve rates of growth of infected and recovered agents. This is captured by eigenvalues, as ever.

CHAPTER 28 : SIGNALS & TIME SERIES MODELS

DISCRETE-TIME DYNAMICS are especially well-suited to the domain of digital signals. Signal processing is itself a substantial field, full of internal principles and perspectives. Nevertheless, there is a role for dynamical systems to play. This chapter is an idiosyncratic introduction to the linear and affine dynamics of digital signal processing. The interested student should take a real signal processing course for a complete picture. The following quick gloss is selected to resonate with the perspectives and applications of previous chapters: there is so much more.

SLIDING WINDOWS

In what follows, a discrete-time sequence $u = u(n)$ will be rechristened a *signal*, to be best imagined as a digital audio signal for most of the chapter. The problem at hand is how to modify this signal for desired effect. For example, in audio signal processing, one might employ a band-pass filter to remove sounds of certain frequency ranges. Various noise-reduction and other techniques are likewise executed by sweeping through the signal, making local modifications. Temporal signals are by no means confined to audio applications: most economic and price data are presented as a time-series in discrete time.

Signal manipulation or processing can be interpreted in terms of discrete time affine dynamical systems, converting the input signal $u = u(n)$ to an output signal $x = x(n)$ based on u .

EXAMPLE: A simple *smoothing filter* converts an input signal u via the equation:

$$Ex = \alpha x + (1 - \alpha)u \quad \Rightarrow \quad x_{n+1} = \alpha x_n + (1 - \alpha)u_n.$$

When the smoothing parameter α satisfies $0 < \alpha < 1$, the natural dynamics are stable: dropping the input signal to zero causes the output to exponentially decay. It is a helpful exercise to simulate what happens in the case of a nonzero input signal u : one sees that increasing α makes the output x into a smoother version of u .

In the previous example, the smoothing is mild, as it is based on a first-order system that can see the input one step behind. For more control, one can use an affine recurrence relation of the form

$$x_{n+1} = \sum_{i=1}^k c_i x_{n+1-i} + u_n.$$

This sweeps through the input signal u with a window of width k to use, using a linear combination of the input and the last k time steps of x to produce the next entry. Such can be converted to a k^{st} -order affine system on the vector

$$\mathbf{x}(n) = \begin{pmatrix} x_n \\ x_{n-1} \\ x_{n-2} \\ \vdots \\ x_{n-k+1} \end{pmatrix}.$$

This is a bit out of sync with the way that discrete-time vectors are usually represented, but the indexing makes sense in the context of signal processing via a sliding window.

ECHO ECHO ECHO

An example will clarify this setup. Consider the humble *echoplex* effect, familiar to electric guitar players. One feeds an audio signal u into the effector, and the output x has an echo applied. There are two parameters to an echo: the gain $\alpha > 0$ (that determines the strength of the echo) and the delay k (that determines

how distant in time the echo appears). This effector is encoded as a recurrence relation of order k :

$$x_{n+1} = \alpha x_{n-k+1} + u_n \quad \Leftrightarrow \quad (E^k - \alpha I)x = E^{k-1}u$$

This inhomogeneous linear system has characteristic polynomial $\lambda^k - \alpha$ with roots equal to $\sqrt[k]{\alpha}$. There is one real positive root; all the rest are evenly-spaced on the circle of radius $\sqrt[k]{\alpha}$ in the complex plane. The gain α has a direct impact on stability of the echo effect. For $0 < \alpha < 1$, a single echo dies out over time, with exponential decay in amplitude; if $\alpha > 1$, the echo increases in strength, leading to overwhelming feedback.

REVERB

When listening to sounds indoors – think of a concert hall or (less pleasant) an auditorium for an early-morning class – the audio tends to reverberate with a shimmer, depending on the geometry of the room. In such settings, one often sees sound treatment on the walls to reduce echo and control reverb for a natural and less distracting sound. Sounds bounce about a room, arriving as a series of staccato echoes that occur so closely together that they blend into a continuous sound. A simple model for reverb comes from the recurrence relation:

$$x_{n+1} = \sum_{i=1}^k \alpha^i x_{n-i+1} + u_n.$$

As before, the gain α is a constant and should satisfy $0 < \alpha < 1$ to avoid instabilities. Converting this to matrix form is revealing:

$$\begin{pmatrix} x_{n+1} \\ x_n \\ x_{n-1} \\ \vdots \\ x_{n-k} \end{pmatrix} = \begin{bmatrix} \alpha & \alpha^2 & \alpha^3 & \dots & \alpha^k \\ 1 & 0 & 0 & \dots & 0 \\ 0 & 1 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & \dots & 1 & 0 \end{bmatrix} \begin{pmatrix} x_n \\ x_{n-1} \\ x_{n-2} \\ \vdots \\ x_{n-k+1} \end{pmatrix} + \begin{pmatrix} u_n \\ 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix}.$$

This should ring a (reverberating) bell, as it is a population model from Chapter 26, with a few unique features: 1) reproduction rates decrease geometrically over cohorts; and 2) there is an influx of “babies” – the input signal comes in at the first cohort.

The characteristic shimmer associated with a reverb arises from the characteristic polynomial of this matrix:

$$\lambda^k - \sum_{i=1}^k \alpha^i \lambda^{k-i} = 0.$$

The roots – the eigenvalues of reverberation – are as in the case of the Leslie model. For $0 < \alpha < 1$, there is a dominant eigenvalue $0 < \lambda_* < 1$ that is stable; the dominant eigenvector

THE RAISING OF HOGS REDUX

In Chapter 13, a discrete-time supply-demand model was considered that had a time delay built in (the cartoon motivation being that it takes two years to bring hogs to market). The resulting 2nd-order recurrence relation was a function of the two-time-period delay. Other settings merit a longer time delay: the markets for semi-conductors and Christmas trees typically involve at least a five-year delay from planning to sales; the market for other products (e.g., lumber) is significantly longer-term.

An identical balance between linear supply and demand curves and a longer k -step time delay leads to a higher order recurrence relation of the form

$$S_0 + aP_n = D_0 - bP_{n+k} \Rightarrow P_{n+k} + \frac{a}{b}P_n = \frac{D_0 - S_0}{b}.$$

There is a unique equilibrium at the market equilibrium price $(D_0 - S_0)/(a + b)$ whose characteristic equation,

$$\lambda^k + \frac{a}{b} = 0 \Rightarrow \lambda = \left(-\frac{a}{b}\right)^{1/k},$$

has k roots equidistributed along the circle of radius $\sqrt[k]{a/b}$ in the complex plane. The general solution (in local coordinates about the equilibrium) will be a linear combination of basis solutions having sine and cosine terms with amplitudes that are exponential functions depending on this radius: the modulus of the eigenvalues. Thus, as in the case $k = 2$, the ratio of supply price sensitivity a to demand price sensitivity b is the key to whether the oscillations about the price equilibrium die out or are amplified over time.

Why is this chapter on signal processing ending in a reverberation of prices and time delays? Note what happens when we write this system out in matrix form, shifting and stacking the prices $(P_n, P_{n-1}, P_{n-2}, \dots, P_{n-k+1})^T$ into a vector? The associated k^{th} order affine system has the form

$$\begin{pmatrix} P_{n+1} \\ P_n \\ P_{n-1} \\ \vdots \\ P_{n-k} \end{pmatrix} = \begin{bmatrix} 0 & 0 & 0 & \cdots & -a/b \\ 1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 & 0 \end{bmatrix} \begin{pmatrix} P_n \\ P_{n-1} \\ P_{n-2} \\ \vdots \\ P_{n-k+1} \end{pmatrix}.$$

This is an example of a *price echo*. The constant input signal $(D_0 - S_0)/b$ is echoed over time with a (negative!) gain controlled by the ratio a/b . Audio echoes and time-delay price aberrations are fundamentally the same phenomenon: temporal signals acted on recurrently. The window size (the number of time-steps from supply decisions to demand discovery) controls the delay of the price echo. The complex eigenvalues induce oscillations in price, either damped or spiraling out of control, depending on the gain.

THE SIGNAL

As in the previous chapter, linear (or affine) dynamics dominate the discussion. This chapter is not focused on methods or applications of an especially *practical* nature, since the audio signal processing here presented is nowhere near the level of sophistication one wants: harmonic analysis, wavelets, information theory, and stochastic methods in general are all crucial in modern signal processing.

The singular point of this chapter is the unifying perspective that dynamical systems thinking permits. What seems like a coincidental resemblance between models here and in previous chapters is in fact a deep connection. Populations and prices, echoes and reverberations – all these are temporal signals subject to autonomous linear dynamics (the primacy of eigenvalues) and nonautonomous inputs.

What would it look like to convert an echo in pricing to reverb? In the delayed economic model of the previous section, one could imagine using an averaging process to spread out the supply/demand disconnect over the time-period (perhaps through annual investment and insurance contracts). This could lead to an affine model of the form

$$S_0 + \sum_{i=1}^k c_i P_{n-i+k} = D_0 - bP_{n+k} ,$$

where the constants c_i depend on the price-averaging model used and completely determine the resulting shimmer in price variations over time. This model can also be updated to deal with changes in baseline supply/demand constants over time. Instead of S_0 and D_0 being fixed, they can (as in practice) vary as functions of time n . Viewed from the perspective of this chapter, the supply-demand disconnect $D_n - S_n$ is the *input signal*, which is acted upon by the model to generate the price signal P_n , with echoes, reverb, or other effects, depending on the model.

A sequence of prices over time contains a muffled history of cycles – boom and bust. Likewise, when looking at population graphs of various nations and peoples over time, one sees the echoes of wartime, peace, prosperity, and decay, which resound down generations, ordered by time. The fundamental unity between these masses of people, spread over decades, and the weekly prices of petrol, are a reminder that time unfolds linearly but on different scales.

CHAPTER 29 : NEUROSCIENCE MODELS

NEUROSCIENCE is an unfathomably rich source of challenging modelling problems and serves as a prime target for applied dynamical systems. The subject is as broad as it is difficult, and this thin chapter can do no justice to the field. The approach employed, with no small risk of idiosyncrasy, will hopscotch along a few historically significant models in anticipation of a different class of contemporary models.

The first choice to be made is that of scale: *local* or *global*? Brains are large, complex, modular systems with intricate dependencies and subsystems. Yet these are all composed of neurons through which electrochemical activity is mediated. Modelling neurons – the local picture – is a logical (and historical) first step. When doing so, it will be important to replicate observed electrical activity, which, as one might guess from experimentally-observed *spike train data*, involves spiking: periodic or recurrent quick elevations in electrical activity, punctuated by inactive cooling-off periods.

FITZHUGH-NAGUMO

One simple and classical model of electrical activity in a neuron is a continuous-time dynamical system on two variables. The first, $V(t)$, is the *membrane potential*; that is, the difference in electrical potential across the neuron boundary. When this excitation potential crosses a certain threshold, it results in an *action potential* signal that neurons use to fire and communicate. The second variable, $W(t)$, is the *recovery* variable, which represents the activation of ion channels that help to reset the membrane potential after an action potential. This variable provides negative feedback to the system. The *Fitzhugh-Nagumo equations* of 1961 are:

$$\frac{d}{dt} \begin{pmatrix} V \\ W \end{pmatrix} = \begin{pmatrix} I + V - W - \frac{1}{3}V^3 \\ \alpha V - \beta W \end{pmatrix},$$

where I is a parameter representing the external current applied to the neuron; α and β are small positive parameters that are tuned to fit experimental data.

This model seems so straightforward – only two dimensions and similar in spirit to the Van der Pol model of Chapter 17 – that one marvels at the wealth of observable features. Given the linear dynamics on W and the cubic-polynomial dynamics on V , one notes that there are anywhere from zero to three equilibria, depending on parameters. Classification of these equilibria reveals sources, sinks, saddles, and spirals depending on the precise values of the parameters. Limit cycles arise frequently: this signifies the biologically relevant phenomena of periodic “*firing*” of the neuron, followed by a recovery/recharge phase. The small values of α, β mean that this is a fast-slow system, with dynamics on V occurring much more quickly than on W (*cf.* spiking).

Other dynamical features appear to be very sensitive to the parameter values chosen. Indeed, it is in the parameter space that the most interesting structures reside. Pitchfork and Hopf bifurcations unsurprisingly occur. However, more delicate and difficult-to-find bifurcations, such as homoclinic bifurcations and saddle-node bifurcations of periodic orbits, occur. There are even some codimension-3 bifurcations in which these others initiate/terminate.

HODGKIN-HUXLEY

The Fitzhugh-Nagumo model is best viewed as a simplistic 2-D reduction of a larger and more biologically faithful model of potentials along an axon. This **Hodgkin-Huxley model** was a grand achievement in modelling in 1952 based on experimental data from a squid giant axon. There are four variables depending on continuous time t . These are:

- ▷ $V(t)$, the voltage potential, as in Fitzhugh-Nagumo;
- ▷ $x_1(t)$, the potassium channel activation probability (in $[0,1]$);
- ▷ $x_2(t)$, the sodium channel activation probability (in $[0,1]$); and
- ▷ $x_3(t)$, the sodium channel *inactivation* probability (in $[0,1]$).

The differential equations are *difficult* to write out in full. In simplified form, one has the following:

$$\begin{aligned}\frac{dV}{dt} &= A_0 + A_1 x_1^4 (V - B_1) + A_2 x_2^3 (V - B_2) + A_3 (V - B_3) \\ \frac{dx_i}{dt} &= \alpha_i(V)(1 - x_i) - \beta_i(V)x_i \quad : \quad i = 1, 2, 3\end{aligned}$$

Here, the A_i and B_i are all constants, but the α_i and β_i are complicated functions of V of the general form

$$\frac{C_0 + C_1 V}{C_2 + e^{C_3 + C_4 V}},$$

where, again, these C_j are constants which are potentially different for each α_i and β_i . One can, for various realistic parameters, find equilibria, limit cycles (indicating neuronal spiking), Hopf bifurcations, and more – much more than in the already complex Fithugh-Nagumo reduction.

This is an extraordinary example of a model that maintains high-fidelity to the experimental systems from which it was derived. This type of dynamical system – with all its variables and parameters – is at once a grand achievement and an admission of defeat. Our decision to begin with the local dynamics has left us little hope to extend to the global setting.

CONNECTING NEURONS

If the difficulties associated with trying to accurately model a single neuron are nontrivial, then whole-brain modelling is enough to test one's courage. A human brain has a bit fewer than 10^{11} neurons, and these are connected in a manner that is much more complex than nearest-neighbor coupling. The connectivity of the brain occurs at multiple scales, impacting specialization and function in ways we are just beginning to understand. What hope can there be for examining a dynamical model?

A small attempt at modelling multiple neurons and their interactions is best begun by keeping the local dynamics as simple as possible. One could begin with a ridiculously reductive model for the *activity* (or *firing rate*) $x(t)$ of a single neuron of the form

$$Dx = -x + b,$$

where $b > 0$ is a constant external input signal for the neuron: the *activation strength*. As readily observed, this has a unique stable equilibrium at $x = b$.

With this intentionally boring local model in place, one can focus on how to link multiple neurons together. Assuming a collection of N neurons, let $x(t)$ be the

time-dependent vector of activities and $\mathbf{b} = (b_i)$ the correspondingly-sized vector of (time-independent) activation strengths with $b_i > 0$ for all i . Repeating the above model with vectors leads once again to an affine system of completely decoupled neurons with a single collective stable equilibrium at $\mathbf{x} = \mathbf{b}$.

To couple the dynamics – to connect neurons – one can use an N -by- N matrix, W of *connection strengths*. The diagonal entries should vanish, $W_{ii} = 0$, since interactions between distinct neurons are encoded in W . There is no need for W to be symmetric: directional influence is biologically realistic. In a so-called *competitive* model, the interactions between neurons are *inhibitory*, meaning that each $W_{ij} \leq 0$ is nonpositive. Still, the resulting system of coupled neurons

$$D\mathbf{x} = -\mathbf{x} + W\mathbf{x} + \mathbf{b} = (W - I)\mathbf{x} + \mathbf{b}$$

is trivial, in that the system is affine with a matrix that has all eigenvalues stable (assuming that $W - I$ is sufficiently nondegenerate), giving rise to a single stable equilibrium.

The reader should ask to what extent this is desirable in a neuroscience model. Although stable equilibria may have an important role in systems of coupled neurons (e.g., as an interpretation for a stored memory), this alone does not suffice to model large-scale behaviors (many stored memories, ability to learn patterns, etc.). A new idea is needed to cross the threshold from affine to fully nonlinear dynamics. This will take the form of an **activation function**, which transforms the $W\mathbf{x} + \mathbf{b}$ term above.

THRESHOLD-LINEAR MODELS

Taking a cue from observations of neural firing, the activation function should create a threshold which suppresses “quiet” neurons but allows more excited neurons to fire. Several types of activation functions have been used in models; the simplest type of nonlinearity that is both effective and biologically relevant is the **ReLU** – the *rectified linear unit*, given by

$$\text{ReLU}(\mathbf{v}) = \max(\mathbf{v}, \mathbf{0}) = \mathbf{v}^+,$$

where the maxima are computed component-wise. Ironically, *ReLU* is a somewhat more familiar function now than in the past, stemming from its frequent use in machine learning and (artificial) neural networks. Incorporating this into the affine model of the previous subsection yields our desideratum: the Threshold Linear Model:

$$D\mathbf{x} = -\mathbf{x} + \text{ReLU}(W\mathbf{x} + \mathbf{b}) = -\mathbf{x} + (W\mathbf{x} + \mathbf{b})^+.$$

(Both functional and superscript notations for the rectified linear unit are given; neither is elegant, and both are found in the literature.) With this model, one can consider the ensuing dynamics as a function of the connection strengths W and the activation strength $\mathbf{b}$. This, still, has a bewildering number of degrees of freedom to explore – how does one begin to systematically classify behaviors?

Recent work of Curto *et al.* has considered the following further simplification. Assume that the activation strengths are all identical (and, without a loss of generality, $b_i = b = 1$). Let each of the connection strengths W_{ij} for $i \neq j$ be of the form $-1 + \epsilon$ or $-1 - \delta$ for some fixed constants $0 < \epsilon < \frac{\delta}{\delta+1}$, where $W_{ij} = -1 - \delta$ indicates that neuron j *strongly inhibits* neuron i , with $W_{ij} = -1 + \epsilon$ meaning that neuron j *weakly inhibits* neuron i . This reduces the connection model to something purely combinatorial – a *directed graph* suffices to encode

this *Combinatorial Threshold-Linear Model*. It can be shown that changes in the values of δ , ϵ , and b do not change behaviors significantly (so long as the bounds on their values are respected).

What types of dynamics are possible in this model? It is both simplistic and unrealistic as compared to the Hodgkin-Huxley model. However, what is lacking in local dynamical complexity is compensated with global structural complexity. For certain directed graphs, it is possible to have systems with multiple stable and unstable equilibria, multiple attracting and repelling limit cycles, and more. The following is one of the many results that have been obtained about the dynamics of combinatorial threshold-linear networks [CTLNs].

LEMMA. If the directed graph associated with a CTLN has no sinks (nodes into which all incident directed edges point into), then there are no stable equilibria and all orbits are bounded.

▷ *Idea:* The negative values of the W_{ij} terms allow for the construction of a trapping region as in Chapter 20. The analysis of the possible fixed points is more involved and uses combinatorial techniques from hyperplane arrangements. ◁

Having all orbits bounded is a type of stability for the global system – one which is very welcome when arguing that this is a model of neuronal dynamics, as infinitely large activity levels would be detrimental to thought. The conclusions of this result remind one of the Poincare-Bendixson Theorem, which, if it held, would imply the presence of stable limit cycles. However, such is not the case: that theorem only holds in 2-D, and CTLNs present systems with no stable equilibria, no stable limit cycles, and no unbounded orbits. What happens in such cases? Simulations lead one's thoughts to wander over images of meandering orbits, never decaying, never repeating, never escaping. A full explanation of this foreshadowing of chaos awaits the next Volume.

CHAPTER 30 : LEARNING & AI MODELS

JUST AS THE BRAIN pulses in periodic bursts, so too do our artificial learners evolve — not smoothly but iteratively, stumbling and correcting, iterating and adjusting. The synapse yields to the weight, the gradient to the will. The tools and perspectives of dynamical systems are relevant to the optimization and control of learning.

FROM NEURONS TO NEURAL NETWORKS

The previous chapter explored models of individual neurons and small networks of coupled neurons. We saw how even simple threshold-linear models could exhibit complex dynamics when arranged in certain network structures. This naturally leads us to consider larger, more structured networks of artificial neurons as models for computation and learning.

The simplest such model is the *perceptron*, introduced by Rosenblatt in 1958. A perceptron takes a vector input $\mathbf{u}$ and produces a scalar output $y \in [0,1]$ via a single-layer affine map followed by an activation function $\sigma: \mathbb{R} \rightarrow \mathbb{R}$, which is often a *sigmoid*. The output, $y = \sigma(\mathbf{w} \cdot \mathbf{u} + b)$, for $\mathbf{w}$ a vector of *weights* and b an offset *bias*, acts as a hyperplane separation classifier: crude but a start.

More salient to modern machine learning is a multi-layer network, where the input vector $\mathbf{u}$ is converted to an output vector $\mathbf{y}$ via a sequence of hidden layers $\mathbf{z}^{(\ell)} = \sigma(\mathbf{W}^{(\ell)} \mathbf{z}^{(\ell-1)} + \mathbf{b}^{(\ell)})$ indexed by layer ℓ , with weight matrices $\mathbf{W}$ and bias vectors $\mathbf{b}$, and activation function σ (acting component-wise). Here, $\ell = 0$ corresponds to the input $\mathbf{u}$, layer N to the output $\mathbf{y}$. For reasons of simple computation, the ReLU function is often used for σ .

Such a layered network yields a (complex nonlinear) function $\mathbf{y} = F(\mathbf{u})$ which, by the Universal Approximation Theorem, can approximate any continuous bounded function for the correct choice of parameters $\Theta = \{\mathbf{W}^{(\ell)}, \mathbf{b}^{(\ell)}\}_\ell$ so long as there is at least one internal hidden layer.

Training a neural net consists of choosing parameters to minimize a given *loss function* $\mathcal{L}$, often defined over a dataset of inputs and corresponding outputs. These parameters – the weights and biases – are the true dynamical variables in machine learning.

GRADIENT LEARNING

Machine learning is not instantaneous: it unfolds in time, step by step. Gradient descent, the workhorse of modern AI, is an inherently discrete-time dynamical system.

Given a model with parameters Θ (weights and biases) rewritten as a vector $\mathbf{x}(n)$ evolving in discrete time; and given a loss function $\mathcal{L}(\mathbf{x})$, gradient descent is:

$$E\mathbf{x} = \mathbf{x} - \eta \nabla \mathcal{L}$$

where $\nabla \mathcal{L}$ is the gradient of the loss function evaluated at $\mathbf{x}$ and $\eta > 0$ is the *learning rate*, a parameter controlling how far the system jumps downhill at each timestep.

Equilibria of this system correspond to critical points of the loss landscape: points where $\nabla \mathcal{L} = \mathbf{0}$. These equilibria are typically:

- ▷ SINKS (desirable, not common),
- ▷ SADDLES (ubiquitous),

▷ SOURCES (rare, dangerous).

Near an equilibrium $\mathbf{x}_*$, one may linearize the map $\mathbf{x} \mapsto \mathbf{x} - \eta \nabla \mathcal{L}$ to obtain the derivative of the right-hand side as $I - \eta H_*$, where H_* is the 2nd derivative (or Hessian) of $\mathcal{L}$ evaluated at the equilibrium. The eigenvalues of $I - \eta H_*$ therefore control stability of learning.

LEMMA: The eigenvalues of $A = I - \eta H_*$ are of the form $\lambda = 1 - \eta\mu$, where μ are the eigenvalues of H_* .

▷ *Proof:* Let $(\mu, \mathbf{v})$ be an eigenvalue-vector pair for H_* , so that $H_*(\mathbf{v}) = \mu\mathbf{v}$. Then $A\mathbf{v} = (I - \eta H_*)\mathbf{v} = (1 - \eta\mu)\mathbf{v}$. ◁

For $\mathbf{x}_*$ a local minimum of $\mathcal{L}$ (a sink), the Hessian H_* is positive definite, and all eigenvalues μ are real and positive. Thus, for sufficiently small learning rate η , the eigenvalues of $(I - \eta H_*)$ are all real and between 0 and 1: a stable equilibrium of discrete-time gradient descent.

GAINING BIFURCATIONS

The learning rate η is a *parameter*, and, as such, can lead to bifurcations. Too small, and convergence is painfully slow; too large, and there can be a loss of stability. Holding the loss function fixed and increasing η , one observes a change in the qualitative behavior of the system.

Revisiting the circumstance above of a local minimum $\mathbf{x}_*$ for a loss function $\mathcal{L}$ with second derivative H_* a positive-definite symmetric matrix, all of whose eigenvalues $\mu_i > 0$ are real and positive, the eigenvalues of the linearization of gradient descent dynamics near $\mathbf{x}_*$ are equal to $\lambda_i = 1 - \eta\mu_i$.

What happens upon increasing the learning parameter η ? At first, learning becomes more efficient, as $\mu_i \rightarrow 0$. However, at some point, some of the eigenvalues become negative, and one might notice alternation in certain coordinates, though with convergence.

For μ_* the largest (dominant, if unique) eigenvalue of H_* , the critical learning rate is $\eta_* = 2/\mu_*$. At this learning rate, the dominant linearized iterative dynamics attains an eigenvalue of $\lambda_* = 1 - \eta_*\mu_* = -1$. This is a classic example of a *period doubling bifurcation*. Increasing the learning rate leads to a loss of stability at the local minimum and the creation of a new, stable, period-2 orbit. Further – unwise – increases in the learning rate can lead to a cascade of period doubling bifurcations and unpredictable dynamics, as foretold in Chapter 9. Empirically, this manifests as training instability: exploding gradients, failure to converge, or chaotic updates.

Beyond bifurcation, there lies opportunity. Oscillations can be tamed via momentum, friction, or rescaling. And in some regimes, controlled instability can help the system escape poor minima or saddles. The trick, as ever, is in the tuning.

DOUBLE DESCENT

The classical narrative of generalization error begins with the bias–variance tradeoff, yielding a U-shaped test error curve as model complexity increases. This story is upended by the phenomenon of *double descent*: as model capacity increases beyond the interpolation threshold—where the model fits the training data exactly—the test error first peaks, then decreases again. Empirically, this

second descent appears in neural networks, kernel machines, and even linear models. Why?

From a dynamical viewpoint, increasing model capacity changes not just what minima exist, but which are accessible under gradient descent. For models just capable of interpolation, the loss landscape is sharp and rugged, and the basin structure of gradient descent becomes fractured. Convergence is sensitive to initialization and noise. But as capacity increases further, the loss landscape flattens again: there are many wide, flat minima, and gradient descent—biased toward such attractors—finds them with ease.

Let $\mathbf{x} \in \mathbb{R}^d$ be the parameter vector of a model, and $\mathcal{L}(\mathbf{x})$ the loss function, with gradient descent yielding the discrete-time dynamical system:

$$E\mathbf{x} = \mathbf{x} - \eta \nabla \mathcal{L}.$$

As model dimension d (and thus capacity) increases, the dynamical behavior of gradient descent shifts profoundly. Initially, near the interpolation threshold, where the model barely fits the training data, the Hessian $H = [D^2\mathcal{L}]$ at or near minima can exhibit a spectrum with some large, isolated eigenvalues $\{\lambda_i\}$. For a given learning rate η , these correspond to directions of high curvature. The component of the error along an eigenvector $\mathbf{v}_i$ scales by $(1 - \eta\lambda_i)$ per step in the linearized dynamics. If λ_i is large, $\eta\lambda_i$ can easily exceed the bifurcation point at which $(1 - \eta\lambda_i) < -1$. This leads to instability (oscillations and divergence) along these sharp directions, making convergence to well-generalizing minima difficult.

However, as d increases significantly into the overparameterized regime, the Hessian's eigenvalue spectrum often undergoes a transformation. Many eigenvalues may collapse towards zero (rendering the Hessian effectively low-rank and indicating numerous flat directions), while the remaining non-zero eigenvalues might become smaller and more uniformly distributed (approaching an isotropic state). This signifies a general flattening of the loss landscape. Consequently, wide, relatively flat *basins of attraction*—corresponding to solutions with smaller Hessian eigenvalues—become prevalent, and the gradient flow within them simplifies.

In this overparameterized regime, a form of implicit regularization emerges directly from the optimization dynamics. Flat minima, which are often correlated with better generalization, become more accessible. While the model has enough capacity to memorize noisy training data (which could correspond to sharp, complex minima), gradient descent navigates this landscape with inherent biases:

- ▷ **STABILITY:** For a fixed learning rate η , directions with large λ_i (sharp minima) risk instability because $\eta\lambda_i > 2$ is more likely. Conversely, directions with small λ_i (flat minima) satisfy the stability condition.
- ▷ **BASIN VOLUME:** Wider, flatter minima often possess larger basins of attraction, increasing the likelihood of random initializations converging to them.

The evolution of gradient descent effectively “filters out” or makes it harder to settle into overly sharp minima, guiding the parameters towards smoother solutions characterized by smaller Hessian eigenvalues. The second descent in test error is not a statistical anomaly but reflects this dynamical phase change in the *accessible* attractor structure of training. The sharpness of the landscape, measured by the magnitude and spread of the Hessian eigenvalues $\{\lambda_i\}$, is

fundamentally altered by overparameterization, which tends to flatten this spectrum.

Thus, double descent is not merely a statistical curiosity—it signals a qualitative shift, akin to a bifurcation, in the dynamical topology of the optimization landscape as explored by gradient descent. Learning, which might have initially struggled to navigate narrow, sharply curved valleys prone to instability, now finds itself more readily drifting into, and settling within, expansive, gently sloping plains that correspond to solutions with better generalization.

GENERATIVE-ADVERSARIAL NETWORKS

Finding minima of the loss function has been our focus thus far. However, a remarkable class of problems in machine learning centers on the saddle point problem:

$$\text{minimize over } \mathbf{x}, \text{ maximize over } \mathbf{y}: \mathcal{L}(\mathbf{x}, \mathbf{y})$$

where $\mathbf{x}$ and $\mathbf{y}$ are sets of parameters controlled by different agents — often called the *generator* and the *discriminator*. A simple example of this arises in producing realistic images of faces: one trains a neural network generator to produce realistic images; along with a neural network discriminator whose job is to determine whether an image is true or generated. These play an adversarial game with true and generated images; the loser in each round has their weights updated, with the goal of converging to a perfect generator-discriminator pair.

Such *Generative Adversarial Networks* (GANs) are common architectures in reinforcement learning and adversarial robustness. When translated into the language of optimization theory, the goal is not to find a minimum of a loss function, but a *saddle*: a point that is minimal in one direction (minimizing errors in detection) and maximal in another (generating passable images).

To study such systems dynamically, consider the following simple case. The generator and discriminator have parameters $\mathbf{x}$ and $\mathbf{y}$, and optimization is with respect to a bilinear loss function $\mathcal{L}(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T A \mathbf{y}$. Under the usual gradient descent, one obtains the following discrete-time system with learning rate η :

$$\mathbf{x}_{n+1} = \mathbf{x}_n - \eta A \mathbf{y}_n \quad : \quad \mathbf{y}_{n+1} = \mathbf{y}_n + \eta A^T \mathbf{x}_n$$

This is a linear dynamical system whose matrix in block format is given by

$$\begin{bmatrix} I & -\eta A \\ \eta A^T & I \end{bmatrix}$$

One can show as an exercise that this matrix has eigenvalues of the form $\lambda_i = 1 \pm i\eta\omega_i$ for certain $\omega_i > 0$ dependent on A ; these are all unstable. Thus, the naïve gradient descent has an eigendecomposition into planes where each plane has an unstable spiral. This is not good.

EXAMPLE: Consider the trivial 2-D example $\mathcal{L} = xy$ with naïve gradient descent given by

$$E \begin{pmatrix} x \\ y \end{pmatrix} = \begin{bmatrix} 1 & -\eta \\ \eta & 1 \end{bmatrix} \begin{pmatrix} x \\ y \end{pmatrix},$$

with eigenvalues $\lambda = 1 \pm i\eta$: clearly unstable, as the magnitude exceed one. To obtain stability, one can use the following *extragradient method*. Let $\tilde{x}_n$ and $\tilde{y}_n$ denote the predictors for x and y at time step n . These are coupled with an interleaving, feeding the predictors into the dynamics as follows:

$$\begin{aligned} \tilde{x}_n &= x_n - \eta y_n & : & \quad \tilde{y}_n = y_n + \eta x_n \\ x_{n+1} &= x_n - \eta \tilde{y}_n & : & \quad y_{n+1} = y_n + \eta \tilde{x}_n \end{aligned}$$

When composed and written as a linear system, one obtains

$$E \begin{pmatrix} x \\ y \end{pmatrix} = \begin{bmatrix} 1 - \eta^2 & -\eta \\ \eta & 1 - \eta^2 \end{bmatrix} \begin{pmatrix} x \\ y \end{pmatrix}.$$

This system has eigenvalues $\lambda = (1 - \eta^2) \pm i\eta$, which unlike the original system, has magnitudes $\sqrt{1 - \eta^2 - \eta^4}$, making the system *stable* for $0 < \eta < 1$.

This example points to a more general approach. To find minimax saddles of a more general loss function $\mathcal{L}(x, y)$, the extragradient method has a coupled predictive step:

$$\begin{pmatrix} \tilde{x} \\ \tilde{y} \end{pmatrix} = \begin{bmatrix} I & -\eta \nabla_x \mathcal{L} \\ \eta \nabla_y \mathcal{L} & I \end{bmatrix} \begin{pmatrix} x \\ y \end{pmatrix} \quad : \quad E \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} x \\ y \end{pmatrix} + \begin{bmatrix} 0 & -\eta \nabla_x \mathcal{L} \\ \eta \nabla_y \mathcal{L} & 0 \end{bmatrix} \begin{pmatrix} \tilde{x} \\ \tilde{y} \end{pmatrix}$$

This anticipates and corrects accordingly. Under suitable assumptions on $\mathcal{L}$, these dynamics will converge to the saddle optimum.

CHAPTER 31 : NETWORKED MODELS

COUPLED & UNCOUPLED dynamical subsystems constitute a thread running through this text, motivating everything from diagonalization to eigendecomposition. In instances where coupling is pairwise, the underlying structure can be modelled as a network of vertices and edges. Dynamics on networks, generalized from neural models, has its own rhythm and flow.

NETWORKS EVERYWHERE

Networks are fundamental structures underpinning many of the complex systems that surround us. In applied dynamical systems, the natural progression from individual dynamics to interconnected multi-agent systems hinges upon understanding how pairwise interactions shape global behavior. Such pairwise couplings can dramatically transform simple dynamics into intricate emergent phenomena, evident in diverse contexts.

EXAMPLES: The following exhibit pairwise coupling.

- ▷ **SOCIAL NETWORKS:** Individuals (nodes) interacting through friendships, professional ties, or digital platforms (edges). Opinion formation, consensus, and polarization arise naturally from peer-to-peer influence.
- ▷ **NEURONAL NETWORKS:** Neurons connected via synapses, where each neuron's electrical activity (firing rate, membrane potential) influences its neighbors. Oscillatory patterns, synchronization, and complex brain rhythms emerge directly from such neural couplings.
- ▷ **ELECTRICAL GRIDS:** Power stations, transformers, and substations (nodes) interconnected by transmission lines (edges). Couplings manifest through voltage and frequency synchronization, with stability or catastrophic failure (blackouts) resulting from changes in the pairwise connections.
- ▷ **ECOLOGICAL NETWORKS:** Species in an ecosystem coupled through predator-prey, competitive, or mutualistic relationships. Dynamics include stable coexistence, cyclic population fluctuations, or extinction cascades stemming directly from pairwise species interactions.
- ▷ **FINANCIAL NETWORKS:** Financial institutions linked by inter-bank loans or trade exposures. Coupling leads to the propagation of risks, contagion effects, and systemic crises through bilateral financial obligations.
- ▷ **EPIDEMIOLOGICAL NETWORKS:** Individuals or communities connected by physical proximity, travel pathways, or communication links. The spread of infectious diseases and information flow are fundamentally shaped by the local pairwise transmission dynamics.

The ubiquity and diversity of such networks motivate the detailed exploration of their common underlying mathematical framework.

THE GRAPH LAPLACIAN

A networked dynamical system begins with local dynamics, supported on the vertices, potentially varying from vertex to vertex. Second is a communication model, whereby neighbors may communicate their states to neighbors and effect influence on the local dynamics. The most fascinating aspects of networked dynamical systems lie in the passage from local to global behaviors – how local vertex-based dynamics coordinate over edges to construe collective systemic behaviors.

Most models for communication on networks come from matrices which encode the graph structure. Let $G = (V, E)$ be an (un)directed graph with vertex set V and edge set E , where E consists of (un)ordered pairs of [distinct] vertices. The following definitions are standard in graph theory:

- ▷ The **adjacency matrix** is the binary matrix A of size $|V| \times |V|$ given by $A_{ij} = 1$ if there is an edge between vertices i and j ; else, the entry is zero. For an undirected graph, this matrix is symmetric.
- ▷ The **incidence matrix** is the matrix B of size $|V| \times |E|$ given by $B_{ij} = 1$ if a directed edge j terminates at vertex i and $B_{ij} = -1$ if a directed edge j originates from vertex i .
- ▷ The **graph Laplacian** is the matrix L of size $|V| \times |V|$ given by the formula

$$L = D - A = BB^T$$

Here, D is the diagonal **degree matrix** whose entry D_{ii} records the number of edges incident to vertex i .

The formula for the graph Laplacian using adjacency and degree matrices is best for computing and intuiting L . It is a wonderful fact that L is well-defined for undirected graphs, and the expression in terms of BB^T is independent of the orientations chosen on the edges. That this is so hints at its deeper roots in algebraic topology and homological algebra, beyond the boundary of this text.

The Laplacian L is especially useful in networked dynamics, as it captures many global properties of the graph. The *nullity* of L (the dimension of the kernel) counts the number of connected components of the graph. Other eigenvalues and eigenvectors of L give more subtle information: this path of inquiry leads to *spectral graph theory*, an absorbing and useful subject. The only result from that subject we will use is the following, the most useful part of which follows directly from the fact that L is a symmetric positive semi-definite matrix.

LEMMA. For any finite graph the eigenvalues of the Laplacian L are all real and non-negative. The number of zero eigenvalues equals the number of connected components of the graph.

COUPLED OSCILLATORS REDUX

Recall the model of coupled spinners from Chapter 5. These models extend readily to networked dynamical systems. Consider an undirected graph $\mathcal{G} = (V, E)$ of “spinners” at V which are connected pairwise across edges E . Each spinner operates at a fixed natural frequency ω and the angle $\theta_v(t)$ of the spinner at vertex $v \in V$ evolves according to:

$$D\theta_v = \omega + \epsilon \sum_{(u,v) \in E} \sin(\theta_u - \theta_v)$$

where the sum is over all edges incident to v . If we consider the synchronized state where all the θ variables are equal, then $\theta_v(t) = \omega t + \theta_0$ for some fixed θ_0 . We could define phases as in Chapter 5; however, this would transform the system to have phase variables on the edges E . Instead, let us focus on what happens at the vertices.

Consider the variables $\psi_v = \theta_v - (\omega t + \theta_0)$ as measuring how far from sync a given vertex is. Then, we have a simpler system of the form

$$D\psi_v = D(\theta_v - (\omega t + \theta_0)) = \epsilon \sum_{(u,v) \in E} \sin(\psi_u - \psi_v),$$

where we can substitute ψ for θ on the right due to the difference. This has an equilibrium at $\Psi = (\psi_v)_v = \mathbf{0}$. The derivative of the right hand side is a matrix

of size $N = \#V$ square. Differentiating yields a $-\epsilon$ at the diagonal (v, v) for each edge attached to v ; and a $+\epsilon$ at (u, v) for each edge between v and u . In the language of graph theory, this is, simply, $-\epsilon L$, for L the Laplacian of $\mathcal{G}$.

Assuming a connected graph, there is a single (dominant) eigenvalue $\lambda_* = 0$, all remaining eigenvalues are strictly negative (since from the Lemma above L has all eigenvalues nonnegative and $\epsilon > 0$). The dominant eigenvector consists of all ones and corresponds to shifting the synchronized solution. For all near-sync initial conditions, the system converges to the synchronized state where $\Psi = \mathbf{0}$.

What is interesting here is not the slow drift back to sync, but rather the appearance of the graph Laplacian and its spectral implications: knowing the eigenvalues determines the rate of convergence back to sync.

EXAMPLE: For a linear graph of vertices v_i for $i = 0 \dots N - 1$ connected sequentially, the eigenvalues of the linearized right hand side are

$$\lambda_j = -2\epsilon \left(1 - \cos\left(\frac{\pi j}{N}\right) \right) \quad : \quad j = 0 \dots N - 1.$$

This has one zero eigenvalue and all the rest negative. The eigenvalue closest to zero (when $j = 1$) determines the speed of convergence.

EXAMPLE: Consider, instead of a linear graph, a cyclic graph of vertices v_i where, in this case, there is an extra edge from v_0 to v_{N-1} . The eigenvalues of the linearized right hand side are subtly different:

$$\lambda_j = -2\epsilon \left(1 - \cos\left(\frac{2\pi j}{N}\right) \right) \quad : \quad j = 0 \dots N - 1.$$

This has one zero eigenvalue and all the rest negative, but with multiplicities for the nonzero eigenvalues. If instead of nearest-neighbor coupling (each spinner is influenced by two neighbors) we have all-to-all coupling – everyone influences everyone else – then the eigenvalues are *very* different: $\lambda_* = 0$ with multiplicity one, and all other eigenvalues are precisely $-\epsilon N$. This implies *much* faster convergence to the sync, especially as N grows, which, in the linear and cyclic cases, leads to dramatically *slower* convergence.

There are additional interesting dynamics that one can consider for coupled oscillators on a general graph. Clearly, disconnected graphs will have decoupled dynamics. The above analyses only look at local dynamics near the synchronized equilibrium: there are other (unstable) equilibria in which certain nodes are in sync and others are completely out-of-sync.

Finally, and most intriguingly, there are sometimes stable periodic orbits for coupled oscillators which are best described as twisted states, in which a wave pulses along a cycle in the graph. Such solutions only exist for graphs with a cycle and the existence and structure of such twisted waves are closely related to the incidence matrix.

DIFFUSION & HEAT ON GRAPHS

The graph Laplacian L plays the rôle of a diffusion operator on networks, analogous to the Laplace–Beltrami operator on a manifold. Let $\mathbf{x} \in \mathbb{R}^n$ represent a scalar field (e.g., temperature) evolving on the vertices of a graph. The *heat equation* takes the form (in continuous or discrete time):

$$D\mathbf{x} = -\alpha L\mathbf{x} \quad : \quad E\mathbf{x} = (I - \alpha L)\mathbf{x} \quad : \quad \alpha > 0$$

Despite differences in formulation, both dynamics converge under mild conditions to equilibria lying in the kernel of the Laplacian:

$$L\mathbf{x} = \mathbf{0} \Rightarrow D\mathbf{x} = \mathbf{0}, \quad E\mathbf{x} = \mathbf{x}$$

These *harmonic states* are characterized by being constant on each connected component of the graph. They represent locally balanced configurations where no further flow occurs—discrete analogues of classical harmonic functions. In both continuous and discrete time, the Laplacian enforces a smoothing effect: differences across edges are suppressed over time, driving the system toward harmonic consensus.

Such Laplacian diffusion appears in many guises: denoising in signal and image processing, computing electrical potentials in resistor networks, modeling passive transport in population dynamics, stabilizing consensus in distributed systems, and constructing geometric representations of high-dimensional data via diffusion maps.

The unifying thread across these domains is this: the graph Laplacian, whether through differential flow or iterative averaging, selects out the “most neutral” or “most balanced” configuration allowed by the topology of the graph. The story of diffusion is thus the story of harmonics—a quiet equilibrium toward which systems inexorably tend.

CONSENSUS OPINION DYNAMICS

The preceding section traced the action of the Laplacian as a smoothing force – a diffusive balm that drives systems toward harmonic consensus. This physical interpretation finds an unexpected but compelling analogue in the social sciences. Where before we modeled heat or electrical potential, now we consider opinions – volatile, polarizing, yet subject to influence, to averaging, to diffusion.

What is an opinion? Is it a proposition held, a value endorsed, a belief entertained, a shade of sentiment? In many models of collective behavior, opinions are treated as *real numbers*, where zero connotes indifference, and positive or negative sentiment comes in degrees. The semantics vary: the number may represent position on a political spectrum, degree of agreement with a given statement, confidence in a forecast, preference for a product, or trust in a source. The real line serves as an ordered gauge of sentiment.

Managing the opinions of multiple agents under mutual influence is reminiscent of the *driver* model introduced in Chapter 5: agents at nodes of a graph update their opinions by blending slightly with those of neighbors. Mathematically, this process is captured by the same Laplacian diffusion equations:

$$D\mathbf{x} = -\alpha L\mathbf{x} \quad : \quad E\mathbf{x} = (I - \alpha L)\mathbf{x} \quad : \quad \alpha > 0$$

Just as before, both continuous and discrete dynamics converge to a harmonic state $\mathbf{x}_\infty$ that is constant on each connected component of the graph. Agreement, it would seem, is the natural consequence of mild mutual influence.

This model dates to the work of Taylor (1968) and DeGroot (1974), who formalized how consensus emerges from repeated local averaging. It is elegant, democratic, and even hopeful. In theory, increasing connectivity—adding edges, integrating communities, expanding access—should accelerate convergence to shared understanding. In theory.

Yet practice suggests otherwise. The modern networked public sphere—wired, watched, and weaponized—rarely trends toward harmony. Instead of consensus, we find contradiction; instead of convergence, cacophony. Why? Real-world dynamics are rarely linear, rarely symmetric, rarely passive. In response, a rich ecology of generalizations has emerged. Some models introduce stubborn agents who resist change. Others allow for antagonism—negative edges—modeling enmity and misinformation. Still others capture polarization, echo chambers,

belief falsification, and strategic deception. What begins as Laplacian diffusion becomes a battleground of coupled nonlinearities. Even so, the Laplacian remains the baseline: the origin point against which deviations are measured. It is the still water from which the ripples of discord are cast.

DISCOURSE SHEAVES

Rather than detail all possible models for opinion dynamics over social networks, we close this chapter with a single flexible model using a novel structure that has independent interest. The following object, called a *discourse sheaf*, is motivated by the idea that in real political or social discourse, there are many ideas being discussed, with a potential for disconnect between public external expressions and personal internal beliefs.

To motivate the construction, imagine a large-scale political or social discourse in which opinions are formulated, shared, and updated across the network. Diverse topics are under discussion, to the extent that it is impossible to identify a single global issue. Indeed, the reader may call to mind their pet issues, noting that these are idiosyncratic, difficult to communicate, and rarely under direct discussion.

Imagine that you – Reader – have a set of core basis beliefs: you have strong opinions on maybe a dozen or two hot-button topics. An Interlocutor with whom you discourse has their own basis beliefs which may differ from yours not only in polarity and intensity, but in essence: each person has a personal basis of core ideals. When in discourse (of human events), the topics of discussion may not intersect directly with either agent’s core beliefs: one often must discuss new ideas and formulate opinions by gut-reaction, these expressions being a function of internal core beliefs.

For a structural model of this scenario, fix an undirected graph $\mathcal{G} = (V, E)$ of agents V who communicate pairwise across edges E . A *discourse sheaf*, $\mathcal{F}$, consists of a “ $\mathcal{G}$ -shaped” collection of vector spaces and linear transformations:

- ▷ For each vertex $v \in V$, let $\mathcal{F}_v$ be a finite-dimensional vector space with an explicit basis. An element $\mathbf{x}_v \in \mathcal{F}_v$ is a *private opinion vector* that registers intensities of personal opinions on the *basis opinions* for that vertex.
- ▷ For each edge $e \in E$, let $\mathcal{F}_e$ be a finite-dimensional vector space with an explicit basis. An element $\mathbf{x}_e \in \mathcal{F}_e$ is a *public discourse vector* that registers intensities of opinions on the set of *basis discourse topics* for that vertex.
- ▷ For each edge e incident to a vertex v let $\mathcal{F}_{v \rightarrow e}: \mathcal{F}_v \rightarrow \mathcal{F}_e$ be the *expression map*, a linear transformation which formulates expression of a public belief vector from private basis opinions.

Modelling the expression maps $\mathcal{F}_{v \rightarrow e}$ as linear transformations is, of course, an assumption. This $\mathcal{F}$ is the state space for a dynamical system that evolves internal beliefs based on interpreting what others say.

EXAMPLE: A discourse sheaf where every $\mathcal{F}_v$ and every $\mathcal{F}_e$ is the same $\mathbb{R}^1$ models a system where every agent and interaction opines on the same topic, as in the previous section. What is interesting is how the maps can amplify, mollify, or even flip expression of opinions – you can selectively lie to your neighbors by inserting a minus sign in the expression map.

There is a far-reaching generalization of the graph Laplacian to such a system $\mathcal{F}$ of vector spaces and linear transformations. The *sheaf* (or *Hodge*) *Laplacian* is a linear transformation acting on the collection of all private opinion vectors:

$$L_{\mathcal{F}} : \bigoplus_{v \in V} \mathcal{F}_v \rightarrow \bigoplus_{v \in V} \mathcal{F}_v$$

This is built from the maps $\mathcal{F}_{v \rightarrow e}$ composed with their transposes in a manner that is best specified with a little more sophisticated language than fits in these margins. The important properties are that $L_{\mathcal{F}}$, like the graph Laplacian that it generalizes, has a nontrivial kernel and is symmetric positive semidefinite, so that all eigenvalues are real and nonnegative.

Given any initial opinion distribution $\mathbf{x}_0 = \mathbf{x}(0)$, Laplacian-based diffusion via

$$D\mathbf{x} = -\alpha L\mathbf{x} \quad : \quad E\mathbf{x} = (I - \alpha L)\mathbf{x} \quad : \quad \alpha > 0$$

converges to a *consensus state* $\mathbf{x}_{\infty}$ in which every agent has expressed agreement with all neighbors. This does not imply that the agents share internal opinions or agree on anything at all: it is an expressed agreement on the delimited topics under discussion. This is perhaps a more realistic model of how political coalitions are formed.

This topic, as with this entire chapter, is a short introduction to a long topic of great interest in current research. The convolution of local dynamics and global network structure is a fruitful topic of inquiry and discussion.