1

INTRODUCTION: WHAT IS COMPLEXITY?

There is all the difference in the world between *knowing about* and *knowing how to do*.
—J. EVANS, *The History and Practice of Ancient Astronomy*, 1997)

1.1 Complexity Is Not Simple

If turbulence is the graveyard of theories, then complexity is surely the tombstone of definitions. Many books on complexity have been written, and the bravest of their authors have attempted to define complexity, with limited success. Being nowhere as courageous, I have simply decided not to try. Although complexity is the central topic of this book, I hereby pledge to steer clear of any attempt to formally define it.

This difficulty in formally defining complexity is actually surprising because we each have our own intuitive definition of what is "complex" and what is not, and we can usually decide pretty quickly if it is one or the other. To most people a Bartok string quartet "sounds" complex, and a drawing by Escher "looks" complex. Such intuitive definitions can even take an egocentric flavor, i.e., an Escher drawing is complex because "I could not draw it" or a Mozart piano piece is complex because "I could not play it."

The many guises of the complex systems to be encountered further in this book often involve many (relatively) simple individual elements interacting locally with one another. This characterization—it should definitely not be considered a definition—does capture a surprisingly wide range of events, structures, or phenomena occurring in the natural world, that most of us would intuitively label as complex. It even applies to many artificial constructs and products of the human mind. While novels by Thomas Pynchon are typically replete with oddball characters, events therein are for the most part constrained by the laws of physics and usually follow a relatively straightforward timeline. What makes Pynchon's novels complex is that they involve many, many such characters interacting with one another. The complexity arises not from the characters themselves, however singularly they may behave, but rather from their mutual interactions over time. Likewise, many of Escher's celebrated drawings[1] are based on the tiling of relatively simple pictorial elements, which undergo slow, gradual change across the drawing. The complexity lies in the higher-level patterns that arise globally from the mutual relationship of neighboring pictorial units, which are themselves (relatively) simple.

Nice and fine perhaps, but turning this into a formal definition of complexity remains an open challenge. One can turn the problem on its head by coming up instead with a definition of what is *not* complex, i.e., a formal definition of "simple." Again purely intuitive and/or egocentric definitions are possible, such as "simple = my five-year-old could do this." Like complexity, simplicity is, to a good part, in the eye of the beholder. I am a physicist by training and an astrophysicist and teacher by trade; I am well aware that my own personal definition of what is "simple" does not intersect fully with that of most people I know. Yet such

[1] See www.mcescher.com/gallery/transformation-prints for reproductions of artwork by Maurits Cornelis Escher.

divergences of opinions are often grounded in the language used to describe and characterize a phenomenon.

Consider, for example, the game of pool.[2] Even without any formal knowledge of energy and momentum conservation, a beginner develops fairly rapidly a good intuitive feel for *how* the cue ball *should* be hit to propel a targeted, numbered ball into a nearby pocket; reliably executing the operation is what requires skill and practice. Now, armed with Newton's laws of motion, and knowing the positions of the pocket and the two participating balls, the required impact point of the cue ball can be *calculated* to arbitrarily high accuracy; the practical problem posed by the production of the proper trajectory of the cue ball, of course, remains. Whichever way one looks at it, the collision of two (perfectly spherical) pool balls is definitely simple, provided it takes place on a perfectly flat table.

If physical laws allow, in principle, the computation of the exact trajectories of two colliding pool balls, the same laws applied repeatedly should also allow generalization to many balls colliding in turn with one another. Experience shows that the situation rapidly degrades as the number of balls increases. I have not played pool much, but that is still enough to state confidently that upon starting the game, no single pool break is ever exactly alike another, despite the fact that the initial configuration of the 15 numbered balls (the "rack") is always the same and geometrically regular—closely packed in a triangular shape. The unfolding of the break depends not just on the speed, trajectory angle, and impact position of the cue ball, but also on the exact distances between the balls in the rack and on whether one ball actually touches another, i.e., on the *exact* position of each ball. For all practical purposes, the break is unpredictable because it exhibits

[2]The reader unfamiliar with this game will find, on the following Wikipedia page, just enough information to make sense of the foregoing discussion: en.wikipedia.org/wiki/Eight-ball.

extreme sensitivity to the initial conditions, even though the interaction between any pair of colliding balls is simple and fully deterministic.

Is complexity then just a matter of sheer number? If the definition of complexity is hiding somewhere in the interactions between many basic elements, then at least from a modeling point of view we may perhaps be in business. If the underlying physical laws are known, computers nowadays allow us to *simulate* the evolution of systems made up of many, many components, to a degree of accuracy presumably limited only by the number of significant digits with which numbers are encoded in the computer's memory. This "brute-force" approach, as straightforward as it may appear in principle, is plagued by many problems, some purely practical but others more fundamental. Looking into these will prove useful to start better pinning down what complexity is not.

1.2 Randomness Is Not Complexity

If we are to seriously consider the brute-force approach to the modeling of complex systems, we first need to get a better feel for what is meant by "large number." One simple (!) example should suffice to quantify this important point.

Consider a medium-size classroom, say a 3 m high room with a 10×10 m floor. With air density at $\rho = 1.225\,\mathrm{kg\,m^{-3}}$, this $300\,\mathrm{m}^3$ volume contains 367 kg of N_2 and O_2 molecules, adding up to some 10^{28} individual molecules. Written out long that number is

$$10{,}000{,}000{,}000{,}000{,}000{,}000{,}000{,}000{,}000.$$

It doesn't look so bad, but this is actually a *very* large number, even by astronomical standards; just consider that the total number of stars in all galaxies within the visible universe is estimated to be in the range 10^{22}–10^{24}. Another way to appreciate the sheer numerical magnitude of 10^{28} is to reflect upon the fact that 10^{28} close-packed sand grains of diameter 0.25 mm—"medium-grade sand" according to the ISO 14688 standard, but quality beach stuff nonetheless—would cover the whole surface of the Earth, oceans included, with a sandy layer 1 km thick. That is how many molecules we need to track—positions and velocities—to "simulate" air in our classroom.

At this writing, the supercomputers with the largest memory can hold up to $\sim 10^3$ TB $= 10^{15}$ bytes in RAM. Assuming 64-bit encoding of position and velocity components, each molecule requires 48 bytes, so that at most 2×10^{13} molecules can be followed "in-RAM."[3] This is equivalent to a cubic volume element of air smaller than a grain of very fine sand. We are a long way from simulating air in our classroom, and let's not even think about weather forecasting! This is a frustrating situation: we know the physical laws governing the motion and interaction of air molecules, but don't have the computing power needed to apply them to our problem.

Now, back to reality. No one in their right mind would seriously advocate such a brute-force approach to atmospheric modeling, even if it were technically possible, and not only because brute force is seldom the optimal modeling strategy. Simply put, complete detailed knowledge of the state of motion of every single air molecule in our classroom is just not useful in practice. When I walk into a classroom, I am typically interested in global measures such as temperature, humidity level, and perhaps the concentration gradient of Magnum 45 aftershave, so as to pinpoint the location of the source and expel the offending emitter.

It is indeed possible to describe, understand, and predict the behavior of gas mixtures, such as air, through the statistical

[3]Molecules also have so-called "internal" degrees of freedom, associated with vibrational and rotational excitation, but for the sake of the present argument these complications can be safely ignored.

definition of global measures based on the physical properties of individual molecules and of the various forces governing their interactions. This statistical approach stands as one of the great successes of nineteenth-century physics. Once again, a simple example can illustrate this point.

The two panels atop figure 1.1 display two different realizations of the spatially random distribution of $N = 300$ particles within the unit square. Even though the horizontal and vertical coordinates of each particle are randomly drawn from a uniform distribution in the unit interval, the resulting spatial distributions are not spatially homogeneous, showing instead clumps and holes, which is expected, considering the relatively small number of particles involved. Viewing these two distributions from a distance, the general look is the same, but comparing closely the two distributions differ completely in detail—not one single red particle on the left is at *exactly* the same position as any single green particle on the right.

Consider now the following procedure: from the center of each unit square, draw a series of concentric circles with increasing radii r; the *particle number density* (ρ, in units of particles per unit area) can be computed by counting the number of particles within each such circle, and dividing by its surface area πr^2. Mathematically, this would be written as

$$\rho(r) = \frac{1}{\pi r^2} \sum_{n=1}^{N} \begin{cases} 1 & \text{if } x_n^2 + y_n^2 \le r^2, \\ 0 & \text{otherwise.} \end{cases} \tag{1.1}$$

Clearly, as the radius r is made larger, more and more particles are contained within the corresponding circles, making the sum in equation (1.1) larger, but the area πr^2 also increases, so it is not entirely clear a priori how the density will vary as the radius r increases. The bottom panel of figure 1.1 shows the results of this exercise, applied now to two realizations of not 300 but $N = 10^7$ particles, again randomly distributed in the unit

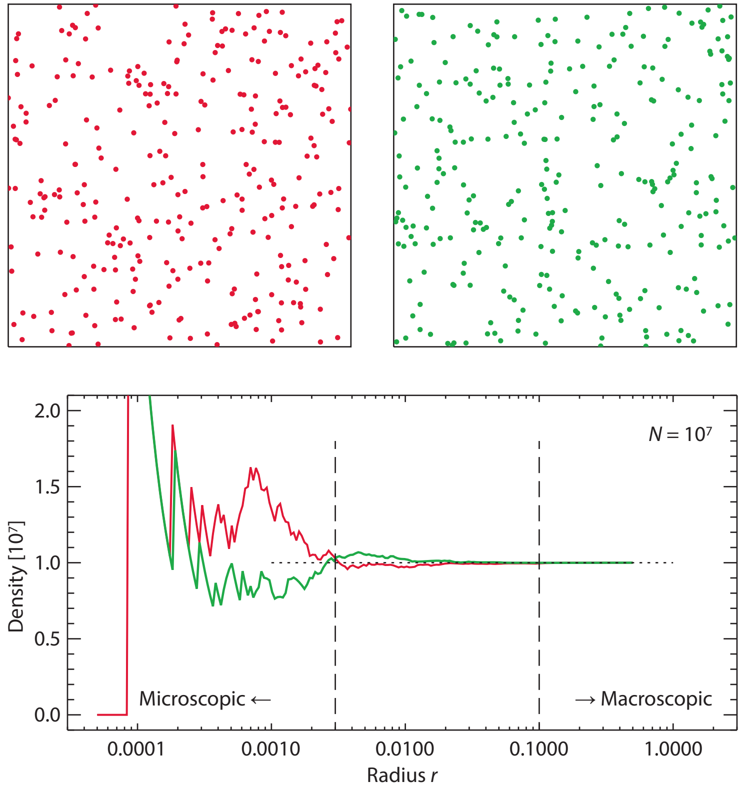

Figure 1.1. Going from the microscopic to the macroscopic scale. The top panels show two distinct random distributions of $N = 300$ particles in the unit square. The bottom panel shows the result of using equation (1.1) to calculate the particle density, based on a series of circles of increasing radii, concentric and centered on the middle of the unit square, now for two distinct random distributions of $N = 10^7$ particles. Note the logarithmic horizontal axis. The resulting density curves differ completely for radii smaller than a few times the mean interparticle distance $\delta = 0.0003$, but converge to the expected value of 10^7 particles per unit area for radii much larger than this distance (see text).

square. The statistically uniform packing of $N = 10^7$ particles in the unit square implies a typical interparticle distance of order $\delta \simeq 1/\sqrt{N} \sim 0.0003$ here. For radii r of this order or smaller, in equation (1.1) the computed density value is critically dependent on the exact position of individual particles, and for $r < \delta$ is it quite possible that no particle is contained within the circle, leading to $\rho = 0$. This is what is happening for the red curve in figure 1.1 up to $r \simeq 0.0001$, while in the case of the distribution associated with the green curve it just so happens that a clump of particles is located at the center of the unit square, leading to abnormally large values for the density even for radii smaller than δ. Nonetheless, as r becomes much larger than δ, both curves converge to the expected value $\rho = 10^7$ particles per unit area.

Figure 1.1 illustrates a feature that will be encountered repeatedly in subsequent chapters of this book, namely, *scale separation*. At the *microscopic* scale (looking at the top panels of figure 1.1 up close) individual particles can be distinguished, and the description of the system requires the specification of their positions, and eventually their velocities and internal states, if any. In contrast, at the *macroscopic scale* (looking at the top panels of figure 1.1 from far back), global properties can be defined that are independent of details at the microscopic scale. Of course, if two systems are strictly identical at the microscopic level, their global properties will also be the same. What is more interesting is when two systems differ at the microscopic level, such as in the two top panels of figure 1.1, but have the same statistical properties (here, x and y coordinates are uniformly distributed in the unit interval); then their physical properties at the macroscopic scale, such as density, will also be the same.

It is worth reflecting a bit more upon this whole argument in order to fully appreciate under which conditions global properties such as density can be meaningfully defined. Considering the statistical nature of the system, one may be tempted to conclude that what matters most is that N be large; but what do we mean by "large"? Large with respect to what? The crux is really that a good separation of scale should exist between the microscopic and macroscopic. The interparticle distance δ (setting the microscopic scale) must be much smaller than the macroscopic scale L at which global properties are defined; in other words, N should be large enough so that $\delta \ll L$. The two vertical dashed lines in figure 1.1 have been drawn to indicate the scale boundaries of the microscopic and macroscopic regimes; the exact values of r chosen are somewhat arbitrary, but a good separation of scale implies that these two boundaries should be as far as possible from one another. In the case of the air in our hypothetical classroom, $\delta \simeq 3 \times 10^{-9}$ m, so that with a macroscopic length scale ~ 1 m, scale separation is very well satisfied.

What happens in the intermediate scale regime, i.e., between the two dashed lines in figure 1.1, is an extremely interesting question. Typically, meaningful global properties cannot be defined, and N is too large to be computationally tractable as a direct simulation. In closed thermodynamic systems (such as the air in our classroom), also lurking somewhere in this twilight zone of sorts is the directionality of time: (elastic) collisions between any two molecules are entirely time reversible, but macroscopic behavior, such as the spread of olfactorily unpleasant aftershave molecules from their source, is not, even though it ultimately arises from time-reversible collisions. Fascinating as this may be, it is a different story, so we should return to complexity since this is complex enough already.

If large N and scale separation are necessary conditions for the meaningful definition of macroscopic variables, they are not sufficient conditions. In generating the two top panels of figure 1.1, particles are added one by one by drawing random numbers from the unit interval to set their horizontal (x) and vertical coordinates (y). The generation of the (x, y) coordinates for a given particle is entirely independent of the positions of the particles already placed in the unit square; particle positions are entirely *uncorrelated*. In subsequent chapters, we will repeatedly encounter situations where the "addition" of a particle to a system is entirely set by the locations of the particles already in the system. Particle positions are then strongly correlated, and through these correlations complexity can persist at all scales up to the macroscopic.

To sum up the argument, while systems made up of many interacting elements may appear quite complex at their microscopic scale, there are circumstances under which their behavior at the macroscopic scale can be subsumed into a few global quantities for which simple evolutionary rules can be constructed or inferred experimentally. The take-home message here then is the following: although complex natural systems often involve a large number of (relatively) simple individual elements interacting locally with one another, not all systems made up of many interacting elements exhibit complexity in the sense to be developed throughout this book. The 10^{28} air molecules in our model classroom, despite their astronomically large number and ever-occurring collisions with one another, collectively add up to a simple system.

1.3 Chaos Is Not Complexity

Complex behavior can actually be generated in systems of very few interacting elements. Chaotic dynamics is arguably the best known and most fascinating generator of such behavior, and there is no doubt that patterns and structures produced by systems exhibiting chaotic dynamics are "complex," at least in the intuitive sense alluded to earlier.

Practically speaking, generators of chaotic dynamics can be quite simple indeed. The *logistic map*, a very simple model of population growth under limited carrying capacity of the environment, provides an excellent case in point. Consider a biological species with a yearly reproduction cycle, and let x_n measure the population size at year n. Under the logistic model of population growth, the population size at year $n+1$ is given by

$$x_{n+1} = A\,x_n(1 - x_n), \qquad n = 0, 1, 2, \ldots, \tag{1.2}$$

where A is a positive constant, and x_0 is the initial population. Depending on the chosen numerical value of A, the iterative sequence $x_0, x_1, x_2, \ldots$ can converge to zero, or to a fixed value, or oscillate periodically, multiperiodically, or aperiodically as a function of the iteration number n. These behaviors are best visualized by constructing a *bifurcation diagram*, as in the bottom-left panel of figure 1.2. The idea is to plot successive values of x_n produced with a given value of A, excluding, if needed, the transient phase during which the initial value x_0 converges to its final value or set of values, and repeating this process for progressively larger values of A. Here, for values of $1.0 < A < 3.0$, the iterative sequence converges to a fixed nonzero numerical value, which gradually increases with increasing A; this leads to a slanted line in the bifurcation diagram, as successive values of x_n for a given A are all plotted atop one another. Once A exceeds 3.0, the iterates alternate between two values, leading to a split into two branches in the bifurcation diagram. Further increases in A lead to successive splittings of the various branches, until the chaotic regime is reached, at which point the iterate x_n varies aperiodically. This is a classical example of transition to chaos through a period-doubling cascade.

The bifurcation diagram for the logistic map is certainly complex in the vernacular sense of the word; most people would certainly have a hard time drawing it with pencil and paper. There is in fact much more to it than that. The series of nested close-ups in figure 1.2 zooms in on the end point of the period-doubling cascade, on a branch of the primary transition to chaos. No matter the zooming level, the successive bifurcations have the same shape and topology. This *self-similarity* is the hallmark of scale invariance, and marks the bifurcation diagram as a fractal

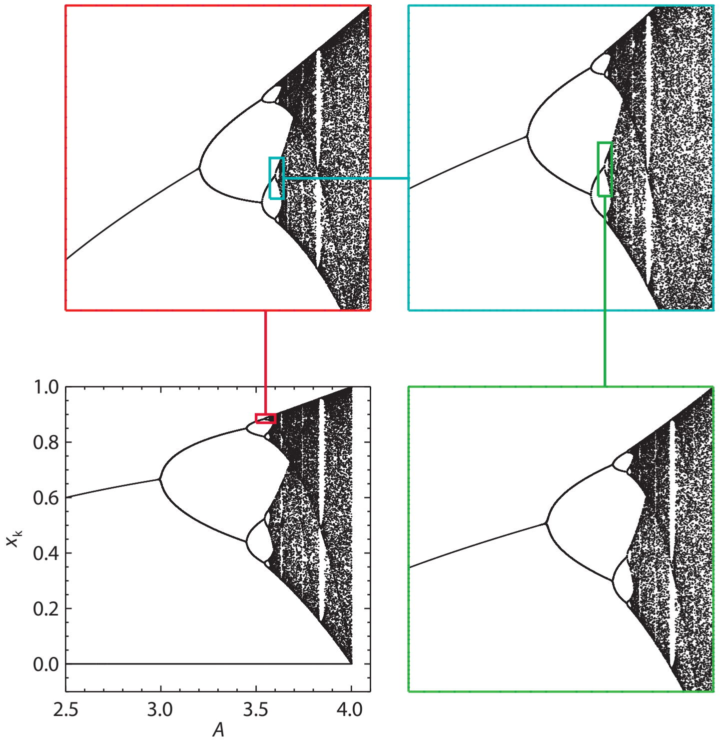

Figure 1.2. Bifurcation diagram for the logistic map (bottom left), as given by equation (1.2). The first bifurcation from the trivial solution $x_n = 0$ occurs at $A = 1.0$, off to the left on the horizontal scale. The other three frames show successive nested close-ups (red→blue→green) on the period-doubling cascade to chaos.

structure. We will have a lot more to say on scale invariance and fractals in subsequent chapters, as these also arise in the many complex systems to be encountered throughout this book.

Chaotic systems such as the logistic map also exhibit *structural sensitivity*, in the sense that they can exhibit qualitative changes of behavior when control parameters—here the numerical constant A—undergo small variations. For example, in the case of the logistic map, increasing A beyond the value 3.0 causes the iterate x_n to alternate between a low and a high value, whereas before, it converged to a single numerical value. In the chaotic regime the map is also characterized by sensitivity to initial conditions, in that the numerical difference between the x_n's of two sequences, differing by an infinitesimally small amount at $n = 0$, is amplified exponentially in subsequent iterations.

Many complex systems to be encountered in the following chapters exhibit similar sensitivities, but for entirely different reasons, usually associated with the existence of long-range correlations established within the system in the course of its prior evolution, through simple and local interactions between their many constitutive elements. In contrast, the cleanest examples of chaotic systems involve a few elements (or degrees of freedom), subject to strong nonlinear coupling. Although such chaotic systems generate patterns and behavior that are complex in the intuitive sense of the word, in and of themselves they are not complex in the sense to be developed in this book.

1.4 Open Dissipative Systems

One common feature of systems generating complexity is that they are *open* and *dissipative*. Pool can serve us well once again in providing a simple example of these notions. After a pool break, the moving balls eventually slow down to rest (with hopefully at least one falling into a pocket in the process). This occurs because of kinetic energy loss due to friction on the table's carpet, and not-quite-elastic collisions with the table's bumpers. The system jointly defined by the moving balls is *closed* because it is subjected to no energy input after the initial break, and is *dissipative* because that energy is slowly lost to friction (and ultimately, heat) until the system reaches its lowest energy equilibrium state: all balls at rest.

Imagine now that the pool table in located inside a ship sailing a rough sea, so that the table is ever slowly and more or less randomly tilted back and forth. Following the break, the balls may slow down to some extent, but will not come to rest since they intermittently pick up energy from the moving table. They will also sometimes temporarily lose kinetic energy of course, for example when finding themselves moving “uphill” due to an unfavorable tilt of the table. But the point is that the balls will not stop moving (well, until they all end up in pockets) no matter how long we wait. A player somehow unaware of the ship’s rock-and-roll would undoubtedly wonder at the curiously curved trajectories and spontaneous acceleration and deceleration of the moving pool balls, and perhaps conclude that their seventh piña colada was one too many.

In this seafaring pool situation, the equilibrium state is one where, on average, the table’s motion injects energy into the system at the same rate as it gets dissipated into heat by friction. The system is still dissipative but is now also *open*, in that it benefits from an input of energy from an external source. At equilibrium, there is as much energy entering the system as is being dissipated, but the equilibrium state is now more interesting: the balls are perpetually moving and colliding, a consequence of energy moving *through* the system.

A most striking property of open dissipative systems is their ability to generate large-scale structures or patterns persisting far longer than the dynamical timescales governing the interactions of microscopic constituents. A waterfall provides a particularly simple example; it persists with its global shape unchanged for times much, much longer that the time taken by an individual water molecule passing through it. As a physical object, the waterfall is obviously “made up” of water molecules, but as a spatiotemporal structure the identity of its individual water molecules is entirely irrelevant. Yet, block off the water supply upstream, and the waterfall disappears on the (short) timescale it takes a water molecule to traverse it. The waterfall persists as a structure only because water flows through it, i.e., the waterfall is an open system.

This line of argument carries over to systems far more intricate than a “simple” waterfall. Consider, for example, Earth’s climate; now that is certainly a complex system in any sense of the word. Climate collects a very wide range of phenomena developing on an equally wide range of spatial and temporal scales: the seasonal cycle, large-scale atmospheric wind patterns such as the jet stream, oceanic currents, recurrent global patterns such as El Niño, tropical storms, and down in scale to thunderstorms and tornadoes, to name but a few. Solar radiative energy entering the atmosphere from above is the energy source ultimately powering all these phenomena. Yet, globally the Earth remains in thermal equilibrium, with as much energy absorbed on the dayside as is radiated back into space over its complete surface in the course of a day. Earth is an open system, with solar energy flowing in and out. If the Sun were to suddenly stop shining, the pole–equator temperature gradient would vanish and all atmospheric and oceanic fluid motions would inexorably grind to a halt, much like the pool balls eventually do after a break on a fixed table. Everything we call climate is just a temporary channeling of a small part of the “input” solar radiative energy absorbed by Earth, all ultimately liberated as heat via viscous dissipation and radiated back into space. The climate maintains its complexity, and generates persistent large-scale weather patterns—the equivalent of our waterfall—by tapping into the energy flowing *through* Earth’s atmosphere, surface, and oceans. Earth is an open dissipative system on a very grand scale.

Most complex systems investigated in this book, although quite simple in comparison to Earth’s climate, are open dissipative systems in the same sense. They benefit from an outside source of energy, and include one or more mechanisms allowing energy to be evacuated at their boundary or to be dissipated internally (or both).

1.5 Natural Complexity

Although I have wriggled away from formally defining complexity, considering the title of this book I do owe it to the reader to at least clarify what I mean by *natural complexity*, and how this relates to complexity in general.

Exquisitely complex phenomena can be produced in the laboratory under well-controlled experimental conditions. In the field of physics alone, phase transitions and fluid instabilities offer a number of truly spectacular examples. In contrast, the systems investigated throughout this book are idealizations of naturally occurring phenomena characterized by the autonomous generation of structures and patterns at macroscopic scales that are not directed or controlled at the macroscopic level or by some agent external to the system, but arise instead "naturally" from dynamical interactions at the microscopic level. This is one mouthful of a characterization, but it does apply to natural phenomena as diverse as avalanches, earthquakes, solar flares, epidemics, and ant colonies, to name but a few.

Each chapter in this book presents a simple (!) computational model of such natural complex phenomena. That natural complexity can be studied using simple computer-based models may read like a compounded contradiction in terms, but in fact it is not, and this relates to another keyword in this book's title: *modeling*. In the sciences we make models—whether in the form of mathematical equations, computer simulations, or laboratory experiments—in order to isolate whatever phenomenon is of interest from secondary "details," so as to facilitate our understanding of the said phenomenon. A good model is seldom one which includes as much detail as possible for the system under study, but is instead one just detailed enough to answer our specific questions regarding the phenomenon of interest. Modeling is thus a bit of an art, and it is entirely legitimate to construct distinct models of the same given phenomenon, each aimed at understanding a distinct aspect.

To many a practicing geologist or epidemiologist, the claim that the very simple computational models developed in the following chapters have anything to do with real earthquakes or real epidemics may well be deemed professionally offensive, or at best dismissed as an infantile nerdy joke. Such reactions are quite natural, considering that, still today, in most hard sciences explanatory frameworks tend to be strongly reductionist, in the sense that explanations of global behaviors are sought and formulated preferentially in terms of laws operating at the microscopic level. My own field of enquiry, physics, has in fact pretty much set the standard for this approach. In contrast, in the many complex systems modeled in this book, great liberty is often taken in replacing the physically correct laws by largely ad hoc rules, more or less loosely inspired by the real thing. In part because of this great simplification at the microscopic level, what these models do manage to capture is the wide separation of scales often inherent in the natural systems or phenomena under consideration. Such models should thus be considered as complementary to conventional approaches based on rigorous ab initio formulation of microscopic laws, which often end up severely limited in the range of scales they can capture.

This apology for simple models is also motivated, albeit indirectly, by my pledge not to formally define complexity. Instead, *you* will have to develop your own intuitive understanding of it, and if along the way you come up with your own convincing formal definition of complexity, all the better! To pick up on the quote opening this introductory chapter, there is all the difference in the world between theory and practice, between knowledge and know-how. This takes us to the final keyword of this book's title: *handbook*. This is a "how-to" book; its practical aim is to provide material and guidance to allow you to learn about complexity through hands-on experimentation with complex systems. This will mean coding and running computer programs, and analyzing and plotting their output.

1.6 About the Computer Programs Listed in This Book

My favorite book on magnetohydrodynamics opens its preface with the statement, "Prefaces are rarely inspiring and, one suspects, seldom read." I very much suspect so as well, and consequently opted to close this introductory chapter with what would conventionally be preface material, to increase the probability that it will actually be read, because it is really important stuff.

If this book is to be a useful learning tool, it is *essential* for the reader to code up and run programs, and modify them to carry out at least some of the additional exercises and computational explorations proposed at the end of each chapter, including at least a few of the Grand Challenges. Having for many years taught introductory computational physics to the first-semester physics cohort at my home institution, I realize full well that this can be quite a tall order for those without prior programming experience, and, at first, a major obstacle to learning. Accordingly, in developing the models and computer codes listed throughout this book, I have opted to retain the same design principle as in the aforementioned introductory class:

1. There are no programming prerequisites; detailed explanations accompany every computer code listed.
2. The code listings for all models introduced in every chapter must fit on one page, sometimes including basic graphics commands (a single exception to this rule does occur, in chapter 10).
3. All computer programs listed use only the most basic coding elements, common to all computing languages: arrays, loops, conditional statements, and functions. Appendix A provides a description of these basic coding elements and their syntax.
4. Computing-language-specific capabilities, including predefined high-level functions, are avoided to the largest extent possible.

5. Clarity and ease of understanding of the codes themselves is given precedence over run-time performance or "coding elegance."

Each chapter provides a complete code listing (including minimal plotting/graphics commands) allowing simulation results presented therein to be reproduced. These are provided in the programming language Python, even though most of the simulation codes introduced throughout this book were originally designed in the C or IDL programming languages. The use of Python is motivated primarily by (1) its availability as free-of-charge, public-domain software, with excellent on-line documentation, (2) the availability of outstanding public-domain plotting and graphics libraries, and (3) its rising "standard" status for university-level teaching. Regarding this latter point, by now I am an old enough monkey to have seen many such pedagogical computing languages rise and fall (how many out there remember BASIC? APL? PASCAL?). However, in view of the third design principle above, the choice of a computing language should be largely irrelevant, as the source codes[4] should be easy to "translate" into any other computing language. This wishful expectation was subjected to a real-life reverse test in the summer of 2015: two physics undergraduates in my department worked their way through an early, C-version of this book, recoding everything in Python. Both had some prior coding experience in C, but not in Python; nonetheless few difficulties were encountered with the translation process.

The above design principles also have significant drawbacks. The simulation codes are usually very suboptimal from the point of view of run-time speed. Readers with programming

[4]Strictly speaking, what I refer to here as "source codes" should be called "scripts," since Python instructions are "interpreted," rather than compiled and executed. While well aware of the distinction, throughout this book I have opted to retain the more familiar descriptor "source code."

experience, or wishing to develop it, will find many hints for more efficient computational implementation in some of the exercises included at the end of each chapter. Moreover, the codes are often not as elegant as they could be from the programming point of view. Experienced programmers will undoubtedly find some have a FORTRAN flavor, but so be it. Likewise, seasoned Python programmers may be shocked by the extremely sparse use of higher-level Python library functions, which in many cases could have greatly shortened the coding and/or increase run-time execution speed. Again, this simply reflects the fact that code portability and clarity have been given precedence.

A more significant, but unfortunately unavoidable, consequence of my self-imposed requirement to keep computational (as well as mathematical and physical) prerequisites to a minimum, is that some fascinating natural complex phenomena had to be excluded from consideration in this book; most notably among these perhaps, is anything related to fluid turbulence or magnetohydrodynamics, but also some specific natural phenomena such as solar flares, geomagnetic substorms, Earth's climate, or the workings of the immune system or the human brain, if we want to think *really* complex. Nonetheless, a reader working diligently through the book and at least some of the suggested computational explorations, should come out well equipped to engage in the study and modeling of these and other fascinating instances of natural complexity.

1.7 Suggested Further Reading

Countless books on complexity have been published in the last quarter century, at all levels of complexity (both mathematically and conceptually speaking!). Among the many available nonmathematical presentations of the topic, the following early best seller still offers a very good and insightful broad introduction to the topic:

Gell-Mann, M., *The Quark and the Jaguar*, W.H. Freeman (1994).

For something at a similar introductory level but covering more recent developments in the field, see, for example,

Mitchell, M., *Complexity: A Guided Tour*, Oxford University Press (2009).

At a more technical level, the following remains a must-read:

Kauffman, S.A., *The Origin of Order*, Oxford University Press (1993).

With regard to natural complexity and the hands-on, computational approach to the topic, I found much inspiration in and learned an awful lot from

Flake, G.W., *The Computational Beauty of Nature*, MIT Press (1998).

Complexity is covered in chapters 15 through 19, but the book is well worth reading cover to cover. In the same vein, the following is a classic not to be missed:

Resnick, M., *Turtles, Termites, and Traffic Jams*, MIT Press (1994).

Statistical physics and thermodynamics is a standard part of the physics cursus. In my department the topic is currently taught using the following textbook:

Reif, F., *Fundamentals of Statistical and Thermal Physics*, reprint, Waveland Press (2009).

For a more modern view of the subject, including aspects related to complexity, I very much like the following book:

Sethna, J. P., *Entropy, Order Parameters, and Complexity*, Oxford University Press (2006)

Good nonmathematical presentations aimed at a broader audience are however far harder to find. Of the few I know, I would recommend chapter 4 in

Gamow, G., *The Great Physicists from Galileo to Einstein* (1961), reprint, Dover (1988).

The literature on chaos and chaotic dynamics is also immense. At the nontechnical level, see, for example,

Gleick, J., *Chaos: Making a New Science*, Viking Books (1987).

For readers fluent in calculus, I would recommend

Mullin, T., ed., *The Nature of Chaos*, Oxford University Press (1993);

Hilborn, R.C., *Chaos and Nonlinear Dynamics*, 2nd ed., Oxford University Press (2000).

The logistic model of population growth is discussed in detail in chapters 5 and 6 of Mullin's book. The functional and structural relationship between chaos and complexity remains a nebulous affair. Those interested in the topic can find food for thought in

Prigogine, I., and Stengers, I., *Order Out of Chaos*, Bantam Books (1984);

Kaneko, K., "Chaos as a source of complexity and diversity in evolution," in *Artificial Life*, ed. C. Langton, MIT Press (1995).

The M.C. Escher foundation maintains a wonderful website, where reproductions of Escher's art can be viewed and enjoyed; see

http://www.mcescher.com/.

Anyone interested in Escher's use of symmetry and transformations should not miss

Schattschneider, D., *Escher: Visions of Symmetry*, 2nd ed., Abrams (2003).

Finally, next time you have a good block of reading time in front of you and are in the mood for a mind-bending journey into complexity in the broadest sense of the word, fasten your seat belts and dive into

Hofstadter, D.R., *Gödel, Escher, Bach*, Basic Books (1979).