

Chapter 1

Introduction: Preliminary Musings

Cellular automata (CA) are, fundamentally, the simplest mathematical representations of a much broader class of complex systems (where, for the moment, “complex system” means any dynamical system that consists of more than a few – typically nonlinearly – interacting parts). As such, CA have proven to be extremely useful idealizations of the dynamical behavior of many real complex systems, including physical fluids, neural networks, molecular dynamical systems, natural ecologies, military command and control networks, and the economy, among many others. Because of their underlying simplicity, CA are also powerful conceptual engines with which to study general pattern formation. They have already provided critical insights into the self-organization of chemical reaction-diffusion systems, crystal growth, pattern formation on seashells, and phase-transition-like phenomena in vehicular traffic flow, to name but a few examples. On a more practical side, CA may provide the basis for extremely powerful encryption algorithms, a subject about which there has recently been much heated debate. There is even some serious speculation that CA may provide the backbone of a radically new discrete fundamental physics (an idea that we will focus our attention on in Chapter 12). While the history of CA can be traced back to early *Systems’ Theory* ([bert68], [jant75]) and the rigorous (but almost impenetrably dry) mathematical analyses conducted primarily by Russian researchers in the 1930s and 40s, their more recent incarnation as simple models of complexity in nature can arguably be traced to a single landmark review paper published by Wolfram in the *Reviews of Modern Physics* in 1983 [wolf83a]. Before embarking in earnest on our journey to explore some remarkable properties of CA in this book, we begin by giving in this introductory chapter a short historical overview and gentle first exposure to a few simple examples.

The fundamental challenge of physics has always been the understanding of the phenomenologically observed complexity in nature using a minimal set of simple principles. This reductionist program has historically—and for obvious reasons—concentrated mostly on studies of comparatively simple systems, deliberately avoiding *more complex descriptions and phenomena*.

Two powerful paradigms for studying natural systems have been provided by *statistical mechanics*, which deals with an infinite number of degrees of freedom and uses central limit theorems for computing desired observables, and *dynamical systems theory*, which deals with systems having only a relatively few degrees of freedom. During the last two decades there has been a veritable explosion of applications of nonlinear dynamics to physical systems. In particular, very simple dissipative deterministic classical systems have been shown to possess remarkable chaotic dynamical properties. Dynamical systems theory provides the tools for studying these systems, *provided* that their response to external probes is constrained to lie in some low-dimensional subspace of phase space. Neither traditional paradigm, however, is capable of dealing adequately with the collective behavior of *complex systems*, which—although consisting of large, often infinite, number of degrees of freedom—can behave in ways such that central limit theorems do not apply. In contrast to our understanding of irregular motions in continuous systems, a comprehensive analytical description of discrete complex systems, constructed out of large numbers of simple components undergoing nonlinear local interactions, is presently out of reach.

1.1 Complex Systems

What are complex systems? While difficult to define without first introducing a few technical concepts, we can see examples of complex systems just about everywhere we look in nature; from the turbulence in fluids to global weather patterns to beautifully intricate galactic structures to the complexity of living organisms. All such systems share at least this one property: they all consist of a large assemblage of interconnected, mutually (and typically nonlinearly) interacting parts. Moreover, their aggregate behavior is *emergent*. That is to say, the properties of the *whole* are not possessed by, nor are they directly derivable from, any of the *parts*—a water molecule is not a vortex, and a neuron is not conscious. A complex system must therefore be understood not just in terms of the set of components out of which it is constructed, but the topology of the interconnections and interactions among those components.

Gases, fluids, crystals, and lasers are all familiar kinds of complex systems from physics. Chemical reactions, in which a large number of molecules conspire to produce new molecules, are also good examples. In biology, there are DNA molecules built up from amino acids, cells built from molecules, and organisms built from cells. On a larger scale, the national and global economies and human culture as a whole are also complex systems, exhibiting their own brand of global cooperative behavior. One of the most far-reaching ideas of this sort is James Lovelock's controversial "*Gaia*" hypothesis, which asserts that the entire earth—molten core, biological ecosystems, atmospheric weather patterns and all—is essentially one huge, complex organism, delicately balanced on the edge-of-chaos [love79].

Perhaps the quintessential example of a complex system is the human brain, which, consisting of something on the order of 10^{10} neurons with 10^3 - 10^4 connec-

tions per neuron, is arguably the most *complex* complex system on this planet. Somehow, the cooperative dynamics of this vast web of “interconnected and mutually interacting parts” manages to produce a coherent and complex enough structure for the brain to be able to investigate its own behavior.

The emerging new sciences of complexity and complex adaptive systems explore the important question of whether (and/or to what extent) does the behavior of the many seemingly disparate complex systems found in nature—from the very small to the very large—stem from the same fundamental core set of universal principles.

References include monographs (Badii and Politi [badii97], Kauffman [kauff93], Holland [holl95], Mainzer [main93], Weisbuch [weis91]), popularizations (Lewin [lewin92], Waldrop [wald93], and Gell-Mann [gell94]), conference proceedings (Cowan [cowan94], Varela [varela92a], Yates [yates87a]) and a series of lecture notes from the Santa Fe Institute ([stein89], [jen90a], [nadel91], [nadel92], [nadel93], and [nadel95]). An eloquent exegesis of the complexity “worldview” is Kauffman’s *Investigations* [kauff00].

1.1.1 *Short History*

Whenever a new field emerges, many different individuals contribute to its development. This is of course also true of complex systems theory, yet three persons stand out as originating and shaping much of the field: *Alan Turing* [turing36], *John von Neumann* [vonN51], and *Stephen Wolfram* [wolf83a]. In addition, *Stanislaw Ulam* [ulam62] and *John Conway* [berk82] each made specific, original contributions of long-lasting fundamental significance. Table 1.1 shows a brief chronology of a few milestone events in the study of complex systems.

Turing, in 1936, published a landmark proof of what has come to be known as the *Halting Theorem*. Turing’s theorem fundamentally limits what one is able to know about the running of a program on a computer by asserting that there is in general no way to know in advance if an arbitrary program will ever stop running. In other words, there is, in general, no quick and dirty short-cut way of predicting an arbitrary program’s outcome; this is an example of what is called *computational irreducibility*. About five decades later, Wolfram suggested that computational irreducibility is actually a property not just of computers, but of many real physical systems as well [wolf85e].

Cellular automata were conceived in 1948 by John von Neumann, whose motivation was in finding a reductionist model for biological evolution [vonN51]. His ambitious scheme was to abstract the set of primitive logical interactions necessary for the evolution of the complex forms of organization essential for life. In a seminal work, completed by Burks, von Neumann followed a suggestion by Ulam to use discrete rather than continuous dynamics and constructed a two-dimensional automaton capable of self-reproduction. Although it obeyed a complicated dynamics and had a rather large state space, this was the first discrete parallel computational model formally shown to be a universal computer (which implies, in turn, that it is also computationally irreducible). Twenty years later, the mathematician John Conway introduced his well-known *Life game*, which remains among the simplest

Year	Researcher	Discovery
1936	Turing	Formalized concept of computability; universal turing machine
1948	von Neumann	Abstracted the logical structure of life; introduced self-reproducing automata as a means towards developing a reductionist biology
1950	Ulam	Proposed need for having more realistic models for the behavior of complex extended systems
1966	Burks	Completed and described von Neumann's work
1967	von Bertalanffy, <i>et al</i>	<i>System theory</i> applied to human systems
1969	Zuse	Introduced concept of "computing spaces," or digital models of mechanics
1970	Conway	Introduced two-dimensional cellular automaton Life rule
1977	Toffoli	Applied cellular automata directly to modeling physical laws
1983	Wolfram	Wrote a landmark review article on properties of cellular automata that effectively legitimized the field as research endeavor for physicists
1984	Cowan, <i>et al</i>	Santa Fe Institute founded, serving as a pre-eminent center for the interdisciplinary study of complex systems
1984	Toffoli, Wolfram	First cellular automata conference held at MIT, Boston
1987	Langton	First artificial life conference held at the Santa Fe Institute
1992	Varela, <i>et al</i>	First European conference on artificial life

Table 1.1 Some landmark historical developments in the study of cellular automata and complex systems.

known models proven to be computational universal [berk82].

Other important historical landmarks include the founding, in 1984, of the Santa Fe Institute,* which is one of the leading interdisciplinary centers for complex systems theory research; the first conference devoted solely to research in cellular automata (which is a prototypical mathematical model of complex systems), organized by Farmer, Toffoli and Wolfram at MIT in 1984 [farmer84]; and the first artificial life conference, organized by Chris Langton at Los Alamos National Laboratory, in 1987 [lang89].

*See the *Santa Fe Institute's* web site at <http://www.santafe.edu>.

1.2 Cellular Automata

Cellular automata (CA) are a class of spatially and temporally discrete, deterministic mathematical systems characterized by local interaction and an inherently parallel form of evolution. First introduced by von Neumann in the early 1950s to act as simple models of biological self-reproduction, CA are prototypical models for complex systems and processes consisting of a large number of identical, simple, locally interacting components. The study of these systems has generated great interest over the years because of their ability to generate a rich spectrum of very complex patterns of behavior out of sets of relatively simple underlying rules. Moreover, they appear to capture many essential features of complex self-organizing cooperative behavior observed in real systems.

Although much of the theoretical work with CA has been confined to mathematics and computer science, there have been numerous applications to physics, biology, chemistry, biochemistry, and geology, among other disciplines. Some specific examples of phenomena that have been modeled by CA include fluid and chemical turbulence ([d'Hum86], [gerh89]), plant growth [linde89] and the dendritic growth of crystals [kess90], ecological theory [pkip92], DNA evolution, the propagation of infectious diseases [segel99], social dynamics [axtel96], forest fires [bak90], and patterns of electrical activity in neural networks [fran92]. CA have also been used as discrete versions of partial differential equations in one or more spatial variables. They have most recently been used to simulate some aspects of military combat (see [wood88] and [ilach97]).

The best sources of information on CA are conference proceedings and collections of papers, such as the one's edited by Boccara [bocc93a], Gutowitz [guto90a], Preston [prest84] and Wolfram ([wolf83b], [wolf94a]). An excellent review of how CA can be used to model physical systems is given by Toffoli and Margolus [marg87].

While there is an enormous variety of particular CA models—each carefully tailored to fit the requirements of a specific system—most CA models usually possess these five generic characteristics:

- **Discrete lattice of cells:** the system substrate consists of a one-, two- or three-dimensional lattice of cells.
- **Homogeneity:** all cells are equivalent.
- **Discrete states:** each cell takes on one of a finite number of possible discrete states.
- **Local interactions:** each cell interacts only with cells that are in its local neighborhood.
- **Discrete dynamics:** at each discrete unit time, each cell updates its current state according to a transition rule taking into account the states of cells in its neighborhood.

There is a deceptive simplicity to these characteristics. In fact, as we shall see repeatedly in this book, such systems are capable of extremely complicated behavior. For example, although obeying local rules and having no intrinsic length scales other than the size of the neighborhoods about each site, CA can generate global patterns with very long-range order and correlation. CA are dynamically rich enough to be seriously considered as alternative mathematical models for a variety of physical systems.

While one is free to think of CA as being nothing more than formal idealizations of partial differential equations, their real power lies in the fact that they represent a large class of *exactly computable models*: since everything is fundamentally discrete, one need never worry about truncations or the slow accumulation of round-off error. Therefore, any dynamical properties observed to be true for such models take on the full strength of theorems [toff77a].

Exact computability in this sense, however, is achieved only at the cost of being able to obtain approximate solutions. Perturbation analysis, for example, is rendered virtually meaningless in this context. It is not surprising that traditional investigatory methodologies are not very well suited to studies of complex systems. Since the behavior of such models can generally be obtained *only* through explicit simulation, the computer becomes the one absolutely indispensable research tool.

1.2.1 CA & Computation

Since CA are totally discrete information processors, a useful analogy almost immediately presents itself between such models and conventional computers. Simply identify the initial state value specification in a CA with the initial *data* of a computer program, and identify the uniquely prescribed rule by which the sites evolve with the *program* itself. Running the program is equivalent to using the CA rule to evolve generations from the initial state, with successive configurations representing intermediate results of the computation. Unlike the serial computations making up the CPU cycles in conventional computers, however, CA process information in a fully parallel manner. Basic theoretical results have been obtained by exploiting this formal equivalence between CA and computers [wolf85e].

von Neumann's construction of a two-dimensional automaton capable of self-reproduction was the first discrete parallel computational model shown formally to be a *universal computer*. Twenty years later, John Conway introduced his (now, well studied) *Life* rule, which remains among the simplest known models proven rigorously to be a computationally universal system (see example # 2 below).

Computational universality is the ability to implement any finite algorithm and thus evaluate any computable function [wolf85e]. This property can be proven either by showing that the given system is formally equivalent to another system known to be a universal computer or by directly identifying the dynamics of intrinsically generated structures with the state transitions in a *Turing Machine* [garey79].

Apart from granting systems the same formal computational power as conventional digital computers, computational universality places a severe restriction on possible theoretical predictions of system behavior [toff77a]: because universal com-

puters all require the same order of magnitude times and data storage capacities to process particular algorithms, no computational method can exist with which a shortcut route is taken to a particular dynamical outcome. In the strongest possible terms, the long time behavior of computationally universal dynamical systems can be obtained only by direct simulation. No general predictive procedure is possible, even in principle. This implies, for example, that for systems such as von Neumann's self-reproducing automaton, there can neither be an analytical expression that exactly describes its asymptotic behavior nor an equation that defines the long-term behavior that itself can be solved in a time less than it would take the system to evolve (modulo a polynomial function of the number of iteration steps necessary for it to reach its final state). All such computationally irreducible systems share the property that their own evolution effectively defines the most efficient simulation of their behavior.

1.2.2 Why Study CA?

There are at least four partially overlapping motivations for studying CA, which we order according to increasing levels of theoretical significance:

- *As powerful computation engines*
- *As discrete dynamical system simulators*
- *As conceptual vehicles for studying pattern formation and complexity*
- *As original models of fundamental physics*

1.2.2.1 CA as Powerful Computation Engines

CA allow very efficient parallel computational implementations to be made of lattice models in physics and thus for a detailed analysis of many concurrent dynamical processes in nature. Indeed, dedicated hardware represents one of the most promising practical applications of CA modeling. With the help of such hardware, many heretofore intractable but technologically important problems such as fluid flow near and around airplane wings, are becoming computationally accessible for the first time. Weisbuch has aptly compared such dedicated CA processors to being the numerical equivalents of wind tunnels [weis91]. More ambitiously, Toffoli and Margolus foresee hardware implementations of CA systems as embodying the concept of *programmable matter* [toff91]; that is to say, a computationally amorphous "machine" that can be "programed" to act as a numerical wind tunnel one moment, or a sea of fermions the next. Their CAM-8 machine, a powerful new descendent of their earlier CAM-5 and CAM-6 CA simulators, comes close to realizing this exciting possibility.[†] The CAM-6, for example, which is still available commercially as a plug-in card for the PC, provides a 256-by-256 two dimensional array in which each site may contain up to four bits of data and the entire array is scanned, updated and displayed sixty times every second. Its speed is comparable in performance

[†]See MIT's *Information Mechanics Group* web page on the CAM-8 at <http://www.im.lcs.mit.edu/cam8/>.

to that of a CRAY-1 for this kind of application. A new generation of CA-based massively parallel computing machines is ushering in a new era in computational physics and applied mathematics.

1.2.2.2 CA as Discrete Dynamical System Simulators

CA allow systematic investigation of complex phenomena by embodying any number of desirable physical properties. Reversible CA,[†] for example, can be used as laboratories for studying the relationship between microscopic rules and macroscopic behavior—exact computability ensuring that the memory of the initial state is retained exactly for arbitrarily long periods of time. The behavior of computationally universal reversible rules, such as Fredkin’s *Billiard Ball* model [fredkin82], are of more than passing interest since they can conceivably avoid the fundamental lower bound on dissipation associated with the reversibility of logic elements in conventional computers [keyes70].

Suitably generalized discrete models have thus far been constructed and studied for dendritic crystal growth [kess90], spatial patterns generated by reaction-diffusion systems (such as the Belousov-Zhabotinsky reaction [nadore83]; see example # 3 below), Ising models [croutz86], self-organization in neural networks [fran92], discrete multiple soliton-structure dynamics [aizawa90], and turbulence in hydrodynamical systems [chenh88]. Discrete models of turbulence are particularly impressive in that they clearly show that very simple finite dynamical implementations of local conservation laws (defined so that the discrete system is computationally universal) are capable of exactly reproducing continuum system behavior on the macroscale.

1.2.2.3 CA as Conceptual Vehicles for Exploring Pattern Formation

CA can be treated as abstract discrete dynamical systems embodying intrinsically interesting, and potentially novel, behavioral features. The central motivation is to abstract the general principles governing self-organizing structure formation. Haken ([haken77],[haken83]) has pioneered the interdisciplinary study of self-organization, coining the term *synergetics* to describe this emerging new discipline. Because of their underlying simplicity, CA offer a powerful theoretical and simulation research tool. Related interests include formal classification of the dynamical behavior of complex systems; a better understanding of the relationship between the dynamics of continuous and discrete systems; and the quantification of complexity—as a system property. All of these questions are particularly accessible through studying the generic behavior of CA systems.

[†]Reversible CA are characterized by the property that each site value has a unique predecessor neighborhood configuration.

1.2.2.4 CA as Original Models of Fundamental Physics

CA allow studies of radically new discrete dynamical approaches to microscopic physics, exploring the possibility that nature locally and digitally processes its own future states. The entire last chapter of this book is devoted to a prolonged discussion of such potentially ground breaking models of physics. Using the fact that computationally universal systems are capable of arbitrarily complicated behavior (in the sense that they can mimic any computation performed by a conventional computer), the idea is to construct fundamentally discrete field theories to compete with existing continuous models. The emphasis in this class of models is emphatically *not* to construct a lattice-gauge-like theory; rather, in the same way as lattice-gas CA successfully reproduce continuous fluid flow despite never having heard of the Navier-Stokes equations, so the hope is to abstract a set of microphysical laws that reproduce known behavior on the macro scale. A number of interesting ideas have recently been explored. Fredkin [fredkin93] has arguably gone to the furthest extreme by asserting that the universe is, at its core, a CA! We will offer a few of our own speculations on this subject in the last chapter.

1.2.3 Example #1: One-dimensional CA

For a one-dimensional CA, the value of the i^{th} cell at time t – denoted by $c_i(t)$ – evolves in time according to a rule F that is a function of $c_i(t)$ and other cells that are within a range r (on the left and right) of $c_i(t)$:

$$c_i(t) = F(c_{i-r}(t-1), c_{i-r+1}(t-1), \dots, c_{i+r-1}(t-1), c_{i+r}(t-1)). \quad (1.1)$$

Since each cell takes on one of k possible values—that is, $c_i(t) \in \{0, 1, 2, \dots, k\}$ —the rule F is completely defined by specifying the value assigned to each of the k^{2r+1} possible $(2r+1)$ -tuple configurations for a given range- r neighborhood:

$c_{i-r}(t-1)$	\dots	$c_i(t-1)$	\dots	$c_{i+r}(t-1)$	\dots	$c_i(t)$
0		0			0	$F(0,0,\dots,0)$
0		0			1	$F(0,0,\dots,1)$
\vdots		\vdots			\vdots	\vdots
k		k			k	$F(k,k,\dots,k)$

Since F itself assigns any of k values to each of the k^{2r+1} possible $(2r+1)$ -tuples, there are a total of $k^{k^{2r+1}}$ possible rules, which is an exponentially increasing function of both k and r . For the simplest case of nearest neighbors (range $r=1$) and $k=2$ ($c_i = 0$ or 1), for example, there are $2^8 = 256$ possible rules. Increasing the number of values each cell can take on to $k=3$ (but keeping the radius at $r=1$) increases the rule-space size to $3^{3^3} \approx 7 \cdot 10^{12}$.

Figure 1.1 shows the time evolution of a nearest-neighbor (radius $r=1$) rule where c is equal to either 0 or 1. The row of eight boxes at the top of the figure shows the explicit rule-set, where – for visual clarity – a box has been arbitrarily colored

black if the value $c=1$ and white if $c=0$. For each combination of three adjacent cells in generation 0, the rule F assigns a particular value to the next-generation center cell of the triplet. Beginning from an initial state (at time=0) consisting of the value zero everywhere except the center site, that is assigned the value 1, F is applied synchronously at each successive time step to each cell of the lattice. Each generation is represented by a row of cells and time is oriented downwards. The first image shows a blowup of the first five generations of the evolution. The second shows 300 generations. The figure illustrates the fact that simple rules can generate considerable complexity.




Fig. 1.1 Example of a one-dimensional CA starting from a single nonzero seed.

The space-time pattern generated from a single nonzero cell by this particular rule has a number of interesting properties. For example, it consists of a curious mixture of ordered behavior along the left-hand-side and what appears to be disordered behavior along the right-hand-side, separated by a corrugated boundary moving towards the left at a "speed" of about $1/4$ cells per "clock" tick. In fact, it can be shown that, despite starting from an obviously non-random initial state and evolving according to a fixed deterministic rule, the temporal sequence of ver-

tical values is completely random. Systems having the ability to deterministically generate randomness from non-random input are called autoplectic systems.

As another example, consider the rule shown at the top of figure 1.2. Its space-time evolution, starting from a random initial state, is shown at the bottom of the figure. Note that this space-time pattern can be described on two different levels: either on the cell-level, by explicitly reading off the values of the individual cells, or on a higher-level by describing it as a sea of particle-like structures superimposed on a periodic background. In fact, following a small initial transient period, temporal sections of this space-time pattern are always of the form "...BBBBPBB...BB...BBBP'BB...BBBP"BBB...", where "B" is a state of the periodic background consisting of repetitions of the sequence "10011011111000" (with spatial period 14 and temporal period 7), and the P's represent "particles." The particle pattern $P = "11111000"$, for example, repeats every four steps while being displaced two cells to the left; the particle $P = "11101011000"$ repeats every ten steps while being displaced two cells to the right.

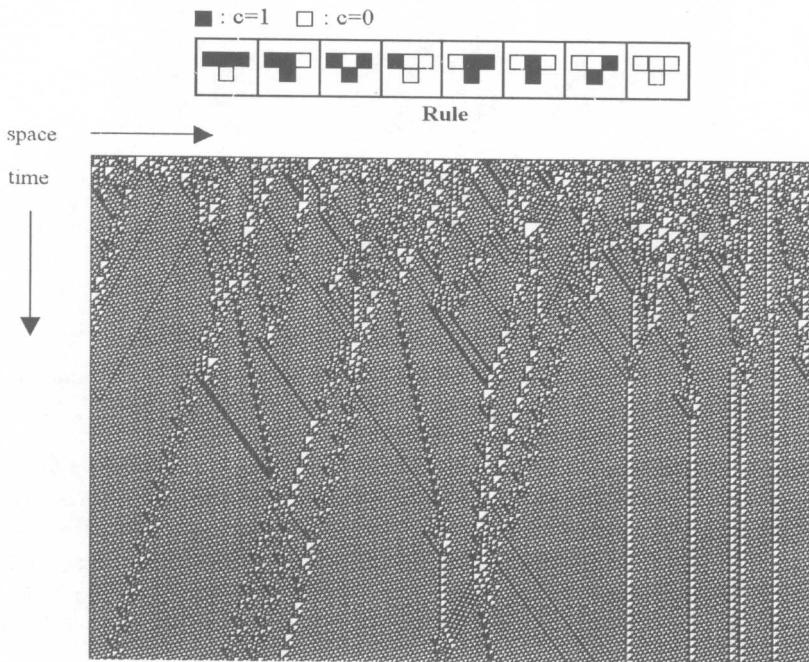


Fig. 1.2 Evolution of a one-dimensional CA starting from a random initial state.

Although the underlying dynamics describing this system is very simple, and entirely deterministic, there is an enormous variety, and complexity, of emergent particle-particle interactions. Such simple systems are powerful reminders that complex higher-level dynamics need not have a complex underlying origin. Indeed, suppose that we had been shown such a space-time pattern but were told nothing whatsoever about its origin. How would we make sense of its dynamics? Perhaps

the only reasonable course of action would be to follow the lead of any good experimental particle-physicist and begin cataloging the various possible particle states and interactions: there are N particles of size s moving to the left with speed v ; when a particle p of type P collides with q of type Q , the result is the set of particles $\{p_1, \dots, p_n\}$; and so on. It would take a tremendous leap of intuition to fathom the utter simplicity of the real dynamics.




Fig. 1.3 Space-time evolution of nine different nearest neighbor one-dimensional CA starting from random initial states.

In general, the behavior of CA is strongly reminiscent of the kinds of behavior observed in continuum dynamical systems, with simple rules yielding steady-state behaviors consisting of fixed points or limit cycles, and complex rules giving rise to behaviors that are analogous to deterministic chaos. In fact, there is extensive empirical evidence suggesting that patterns generated by all (one-dimensional) CA evolving from disordered initial states fall into one of only four basic behavioral classes [wolf83a]:

- Class 1: *Evolution leads to a homogeneous state, in which all cells eventually attain the same value.*
- Class 2: *Evolution leads to either simple stable states or periodic and separated structures.*
- Class 3: *Evolution leads to chaotic nonperiodic patterns.*

- Class 4: Evolution leads to complex, localized propagating structures.

All CA within a given class yield qualitatively similar behavior. While the behaviors of rules belonging to the first three rule classes bear a strong resemblance to those observed in continuous systems – the homogeneous states of class 1 rules, for example, are analogous to fixed-point attracting states in continuous systems, the asymptotically periodic states of class 2 rules are analogous to continuous limit cycles and the chaotic states of class 3 rules are analogous to strange attractors – the more complicated localized structures emerging from class 4 rules do not appear to have any obvious continuous analogues (although such structures are well characterized as being soliton-like in their appearance).

Figure 1.3 shows a few examples of the kinds of space-time patterns generated by binary ($k = 2$) nearest-neighbor ($r = 1$) in one dimension and starting from random initial states.




Fig. 1.4 Examples of CA neighborhoods in two dimensions.

Figure 1.4 shows examples of some commonly used neighborhood structures in two dimensions. These include the von Neumann neighborhood, which consists of the four cells that are horizontally and vertically adjacent to the center cell (left-most graphic appearing in figure 1.4), (2) the Moore neighborhood, that consists of all eight nearest-neighbor cells on a two-dimensional Euclidean lattice (center graphic in figure 1.4), and (3) the Hexagonal neighborhood, that consists of all nearest-neighbor cells on a hexagonal lattice (right-most graphic in figure 1.4).

1.2.4 Example #2: Conway's Life

"Its probable, given a large enough Life space, initially in a random state, that after a long time, intelligent self-reproducing animals will emerge and populate some parts of the space." – John H. Conway [berk82]

Perhaps the most widely known CA is the game of Life, invented by John H. Conway, and popularized extensively by Martin Gardner in his "Mathematical Games" department in *Scientific American* in the early 1970s (see, for example, [gardner70]).

Life is "played" using the 9-neighbor Moore neighborhood (see center graphic in figure 1.4), and consists of (1) seeding a lattice with some pattern of "live" and "dead" cells, and (2) simultaneously (and repeatedly) applying the following three

rules to each cell of the lattice at discrete time steps:

- **Birth:** replace a previously dead cell with a live one if exactly 3 of its neighbors are alive.
- **Death:** replace a previously live cell with a dead one if either (1) the living cell has no more than one live neighbor (i.e. it dies of isolation), or (2) the living cell has more than three neighbors (i.e. it dies of overcrowding).
- **Survival:** retain living cells if they have either 2 or 3 neighbors.

One of the most intriguing patterns in Life is an oscillatory propagating pattern known as the “glider.” Shown on the left-hand-side of figure 1.5, it consists of 5 “live” cells and reproduces itself in a diagonally displaced position once every four iterations.




Fig. 1.5 Glider patterns in Conway's Life.

When the states of Life are projected onto a screen in quick succession by a fast computer, the glider gives the appearance of “walking” across the screen. The propagation of this pseudo-stable structure can also be seen as a self-organized emergent property of the system. The right-hand-side of figure 1.5 shows a still-frame in the evolution of a pattern known as a “glider-gun,” which shoots-out a glider once every 30 iteration steps.

What is remarkable about this very simple appearing rule is that one can show that it is capable of universal computation. This means that with a proper selection of initial conditions (i.e. the initial distribution of “live” and “dead” cells), Life can be turned into a general purpose computer. This fact fundamentally limits the overall predictability of Life's behavior.

The well known *Halting Theorem*, for example, asserts that there cannot exist a general algorithm for predicting when a computer will halt its execution of a given program [gary79]. Given that Life is a universal computer – so that the Halting Theorem applies – this means that one cannot, in general, predict whether a particular starting configuration of live and dead cells will eventually die out. No shortcut is possible, even in principle. The best one can do is to sit back and patiently await Life's own final outcome.

Put another way, this means that if you want to predict Life's long-term behavior with another “model” or by using, say, a partial differential equation, you

are doomed to fail from the outset because its long-term behavior is effectively unpredictable. Life – like all computationally universal systems – defines the most efficient simulation of its own behavior.

1.2.5 Example #3: Belousov–Zhabotinski Reaction

The Belousov-Zhabotinski (BZ) reaction is a chemical reaction consisting of simple organic molecules that is characterized by spectacular oscillating temporal and spatial patterns [zhabo64]. One variant of the BZ reaction involves the reaction of bromate ions with an organic substrate (typically malonic acid) in a sulfuric acid solution with cerium (or some other metal-ion catalyst). When this mixture is allowed to react exothermally at room temperature, interesting temporal and spatial oscillations (i.e. chemical waves) result. The system oscillates, changing from yellow to colorless and back to yellow about twice a minute, with the oscillations typically lasting for over an hour (until the organic substrate is exhausted).

These patterns are an example of what are sometimes called *dissipative structures*, which arise in many complex systems. Dissipative structures are dynamical patterns that retain their organized state by persistently dissipating matter and energy into an otherwise thermodynamically open environment.

Figure 1.6 shows a sample evolution of a CA model of this reaction, in which cells are identified with the reacting molecules, and are colored *black* if they are “active” and *white* if they are “inactive,” according to the reaction rules. The spatial and temporal patterns that emerge from the initially random mixture of states are also a good general example of how CA can be used to model self-organization. The rule shown here is due to Gerhardt and Schuster ([gerh89]; also see color plate 4 in the appendix to chapter three).




Fig. 1.6 Example of self-organization in a two-dimensional CA.

1.2.6 Example #4: Lattice Gases

Lattice gases are micro-level rule-based simulations of macro-level fluid behavior. Lattice-gas models provide a powerful new tool in modeling real fluid behavior ([doolen90], [doolen91]). The idea is to reproduce the desired macroscopic behavior of a fluid by modeling the underlying microscopic dynamics.

It can be shown that three basic ingredients are required to achieve an emergence of a suitable macrodynamics out of a discrete microscopic substrate [hass89]: (1) local thermodynamic equilibrium, (2) conservation laws, and (3) a “scale separation” between the levels at which the microscopic dynamics takes place (among kinetic variables living on a micro-lattice) and the collective motion itself appears (defined by hydrodynamical variable on a macro-lattice). Another critical feature is the symmetry of the underlying lattice.

While there are many variants of the basic model, one can show that there is a well-defined minimal set of rules that define a lattice-gas system whose macroscopic behavior reproduces that predicted by the Navier-Stokes equations exactly.⁵ In other words, there is critical “threshold” of rule size and type that must be met before the continuum fluid behavior is matched, and once that threshold is reached the efficacy of the rule-set is no longer appreciably altered by additional rules respecting the required conservation laws and symmetries.

Figure 1.7 shows a few snapshots of the evolution of a two-dimensional lattice gas starting from an initial condition in which there is a tightly packed region of particles at the center of the lattice. Notice how this central region expands rapidly outward, and is very reminiscent of the effect a dropped stone has on an initially stagnant pool of water. The most striking feature is the circular sound wave, which is circular despite the fact that the microscopic dynamics takes place on a square lattice. The lattice gas rules thus force a symmetry that is not present in the microscopic dynamics to emerge on the macro-scale.




Fig. 1.7 Two-dimensional lattice-gas simulation of a fluid.

1.2.7 Example #5: Collective Behavior in Higher Dimensions

Chate and Manneville [chate92] have examined a wide variety of cellular automata that live in dimensions four, five and higher. They found many interesting rules that while being essentially featureless locally, nonetheless show a remarkably ordered global behavior.

Here again we see a pristine example of the three basic elements of emergence: (1) the global phenomenon (in this case the cell-value density) emerges out of an

⁵The Navier-Stokes equations are the fundamental equations describing incompressible fluid flow; see Chapter 9.

interaction of a large number of simple components (lattice cells of a cellular automaton), (2) there is no evidence of the global phenomenon on the local level, and (3) the global phenomenon obeys a separate dynamics (in this case, quasi-periodicity).

Figure 1.8, for example, plots the probability that a cell has value 1 at time $t+1$ – labeled P_{t+1} – versus the probability that a cell had value 1 at time t – labeled P_t – for a particular four dimensional cellular automaton rule. The rule itself is unimportant, as there are many rules that display essentially the same kind of behavior. The point is that while the behavior of this rule is locally featureless – its space-time diagram would look like noise on a television screen – the global density of cells with value 1 jumps around in quasi-periodic fashion. We emphasize that this quasi-periodicity is a *global* property of the system, and that no evidence for this kind of behavior is apparent in the local dynamics.




Fig. 1.8 Collective behavior of a four dimensional CA (after [chate92]).

1.2.8 Other Variants

There are at least as many variants of the basic CA algorithm as there are ways of generalizing the characteristics of a typical CA system. Here are a few general models:

Asynchronous CA. CA rules are typically defined so that all lattice sites update their values simultaneously throughout the lattice on each time step. A natural generalization is to lift this restriction by allowing *asynchronous* updates [inger84].

Coupled-map Lattices. Another obvious generalization is to lift the restriction that sites can take on only one of a few discrete values. Coupled-map lattices are CA models in which continuity is restored to the state space. That is to say, the cell values are no longer constrained to take on only the values 0 and 1 as in the examples discussed above, but can now take on arbitrary real values. First introduced by Kaneko [kaneko83]-[kaneko93], such systems are simpler than partial differential equations but more complex than generic CA. Coupled-map lattices are discussed in chapter 8.

Probabilistic CA. Probabilistic CA are cellular automata in which the deterministic state-transitions are replaced with specifications of the probabilities of the cell-value assignments. Since such systems have much in common with certain statistical mechanical models, analysis tools from physics are often borrowed for their study. Probabilistic CA are introduced in chapter 8.

Non-homogeneous CA. These are CA in which the state-transition rules are allowed to vary from cell to cell. The simplest such example is one where there are only two different rules randomly distributed throughout the lattice. Kauffman [kauff84] has studied the other extreme in which the lattice is randomly populated with all possible Boolean functions of k inputs.

Mobile CA. These are CA in which some (or all) lattice sites are free to move about the lattice. In effect, mobile CA are primitive models of mobile robots. Typically, their internal state space reflects some features of the local environment within which they are allowed to move and with which they are allowed to interact. An example of mobile CA used to model some aspects of military engagements is discussed in Chapter 12.

Structurally Dynamics CA. Most of the CA that we will encounter throughout this book (indeed, most that are currently being studied!) assume that the underlying lattice remains a passive and static object. The lattice is thus typically an arena for activity, not an active participant in the dynamics. What if the lattice were somehow made an integral part of the dynamics? That is to say, what if the topology -- the sites and connections among sites -- evolved alongside the value states? Structurally dynamic CA are discussed in Chapter 8.

1.3 Outline of Book

The remainder of the book is divided into eleven largely self-contained chapters. Chapter 2 introduces some basic mathematical formalism that will be used throughout the book, including set theory, information theory, graph theory, groups, rings and field theory, and abstract automata. It concludes with a preliminary mathematical discussion of one and two dimensional CA.

Chapter 3 provides a phenomenological introduction to generic CA. The narrative includes both a mathematical description of one, two and three dimensional CA along with guided-tour through a graphics-gallery of typical space-time patterns. Simple ways of parameterizing the space of CA rules are described, and a sketch of the proof of the computational universality of Conway's famous two-dimensional *Life*-rule is also provided.

Chapter 4 covers much of the same ground as chapter 3 but from a more formal dynamical systems theory approach. The discrete CA world is examined in the context of what is known about the behavior of continuous dynamical systems, and a number of important methodological tools developed by dynamical systems theory (i.e. Lyapunov exponents, invariant measures, and various measures of entropy and

dimension) are used to characterize the behavior of simple CA systems.

Chapter 5 provides some examples of purely analytical tools useful for describing CA. It discusses methods of inferring cycle-state structure from global eigenvalue spectra, the enumeration of limit cycles, the use of shift transformations, local structure theory, and Lyapunov functions. Some preliminary research on linking CA behavior with the topological characteristics of the underlying lattice is also described.

Chapter 6 is a short primer on CA and language theory, and provides a basic discussion of formal language theory, the relationship between CA and formal language theory, power spectra of regular languages and reversible computation.

Chapter 7 discusses a variety of topics all of which are related to the class of probabilistic CA (PCA); i.e. CA that involve some elements of probability in their state and/or time-evolution. The chapter begins with a physicist's overview of critical phenomena. Later sections include discussions of the equivalence between PCA and spin models, the critical behavior of PCA, mean-field theory, CA simulation of conventional spin models and a stochastic version of Conway's *Life* rule.

Chapter 8 describes a number of generalized CA models, including reversible CA, coupled-map lattices, quantum CA, reaction-diffusion models, immunologically motivated CA models, random Boolean networks, sandpile models (in the context of self-organized criticality), structurally dynamic CA (in which the temporal evolution of the *value* of individual sites of a lattice are dynamically linked to an evolving lattice *structure*), and simple CA models of combat.

Chapter 9 provides an introductory discussion of a research area that is rapidly growing in importance: *lattice gases*. Lattice gases, which are discretized models of continuous fluids, represent an early success of CA modeling techniques. The chapter begins with a short primer on continuum fluid dynamics and proceeds with a discussion of CA lattice gas models. One of the most important results is the observation that, under certain constraints, the macroscopic behavior of CA models *exactly* reproduces that predicted by the Navier-Stokes equations.

Chapter 10 covers another important field with a great overlap with CA: *neural networks*. Beginning with a short historical survey of what is really an independent field, chapter 10 discusses the Hopfield model, stochastic nets, Boltzman machines, and multi-layered perceptrons.

Chapter 11 contains a very brief introduction to what is rapidly becoming one of the central research areas in complex systems theory; namely, *artificial-life* (AL). Owing its origins to von Neumann's early explorations in self-reproducing automata, AL has blossomed in the last decade from consisting of a few toy worlds hardly more sophisticated than Conway's *Life*-rule universe to intricately rendered 3D artificial universes populated with interacting creatures undergoing an open-ended evolution. Topics include a brief survey of von Neumann's original work, Lindenmeyer systems, Langton's *vants* model, a short primer on genetic algorithms. The chapter concludes with a discussion of an artificial-life-like mobile CA model of land warfare.

The last chapter is a broad survey of a speculative proposition that CA just might one day prove to be even more profound in what they say about how our universe is organized than has heretofore been appreciated. The essence of the

proposition, borrowed from musings by Feynman [fcyn82], Fredkin [fredkin90], Minsky [minsky82], Wheeler [wheel90] and others, is that the universe is fundamentally discrete and obeys, at its core, a simple CA-like dynamics. *Thus begin some final musings and a perhaps a glimpse of a new cosmogony!*

Appendix A provides a brief description of several existing hardware and software tools designed for CA research. Appendix B contains a useful list of CA and more general complexity-related information sources available on the World Wide Web (WWW), subject-sorted into a total of 91 WWW Universal Resource Locator (URL) links in 16 categories. The book is indexed and includes an extensive bibliography.