**Christopher G. Langton**
Complex Systems Group, MS-B213, Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico, 87545, and the Santa Fe Institute, 1120 Canyon Road, Santa Fe, NM 87501

# Life at the Edge of Chaos

In living systems, a dynamics of information has gained control over the dynamics of energy, which determines the behavior of most non-living systems. How has this domestication of the brawn of energy to the will of information come to pass?

By studying the conditions under which a complex dynamics of information can emerge spontaneously in a class of formal systems known as Cellular Automata (CA), we suggest that information can come to dominate the dynamics of physical systems in the vicinity of a second-order (or *critical*) phase transition.

We discuss the implications of this finding for our understanding of the origin and evolution of life and intelligence.

Artificial Life II, SFI Studies in the Sciences of Complexity, vol. X, edited by C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen, Addison-Wesley, 1991 

# 1. INTRODUCTION

What is Life? What is meant by the notion of "the living state?" What properties or characteristics distinguish the living from the non-living?

There have been many criteria proposed to identify the "living state." Biologists today will often point to a long list of properties that are shared by *most* living things. Such lists often include capacities for self-reproduction and metabolism, the exhibition of complex structure and interdependence among parts, reliance on a genetic code, a genotype/phenotype relation, and so forth.[1]

However, the most salient feature that distinguishes living organisms is that their behavior is clearly based on a complex dynamics of information. In living systems, information processing has somehow gained the upper hand over the dynamics of energy that dominates the behavior of most non-living systems.

Therefore, the fundamental question motivating the research described here is:

*Under what conditions can we expect a dynamics of information to emerge spontaneously and come to dominate the behavior of a physical system?*

If we can answer this question we will have gone a long way towards understanding how life could emerge from non-life, not only here on Earth but elsewhere in the Universe as well.

## 1.1 THE DYNAMICS OF CELLULAR AUTOMATA

In this paper, we will address this question via a study of the conditions under which a complex dynamics of information can emerge spontaneously within a class of discrete approximations to physical systems known as cellular automata (CA).

CA are appropriate formalisms within which to pursue this study for a number of reasons.

- CA are spatially extended, nonlinear dynamical systems.
- As nonlinear dynamical systems, CA exhibit the entire spectrum of dynamical behaviors, from fixed-points, through limit cycles, to fully developed chaos.
- CA are capable of supporting universal computation. Thus, they are capable of supporting the most complex known class of information dynamics.
- There is a very general and universal representation scheme for all possible CA: a look-up table. This form of representation allows us to parameterize the space of possible CA, and to search this space in a canonical fashion.
- CA are very physical, a kind of "programmable matter."[45] Thus, what we learn about information dynamics in CA is likely to tell us something about information dynamics in the physical world.

[1]Mayr[35] has collated a representative list of such properties.



So, we can address the fundamental question above in the context of CA by asking:

*Under what conditions can we expect a complex dynamics of information to emerge spontaneously and come to dominate the behavior of a CA?*

## PREVIEW

The investigation of this particular question will take up much of the rest of this paper. First, we will introduce cellular automata more formally, describe methods for sampling the space of possible CA dynamics, and review statistical measures used to characterize these dynamics. Then, we will turn to qualitative and quantitative overviews of CA dynamics, which demonstrate the existence of a *phase transition* in the space of CA. Next, we observe that we can locate the most complex information dynamics in the vicinity of this phase transition, and show how this can serve to explain a great deal about the structure of the space of computations in general. Finally, we return to the main theme of this paper and discuss what all of this has to tell us about life, its origin, and evolution.

# 2. CELLULAR AUTOMATA

Formally, a cellular automaton is a $D$-dimensional lattice with a finite-state automaton (FSA) residing at each lattice site. Each automaton takes as input the states of the automata within some *finite, local* region of the lattice, defined by a neighborhood template $\mathcal{N}$, where the dimension of $\mathcal{N} \leq D$. The size of the neighborhood template, $|\mathcal{N}|$, is just the number of lattice points covered by $\mathcal{N}$. By convention, an automaton is considered to be a member of its own neighborhood.

Each FSA consists of a finite set of *cell states* $\Sigma$, a finite *input alphabet* $\alpha$, and a *transition function* $\Delta$, which is a mapping from the set of neighborhood states to the set of cell states. Letting $N = |\mathcal{N}|$:

$$\Delta : \Sigma^N \rightarrow \Sigma.$$

The *state* of a neighborhood is the cross product of the states of the FSA covered by the neighborhood template. Thus, the input alphabet $\alpha$ for each automaton consists of the set of possible neighborhood states: $\alpha = \Sigma^N$. Letting $K = |\Sigma|$ (the number of cell states) the size of $\alpha$ is equal to the number of possible neighborhood states

$$|\alpha| = |\Delta| = |\Sigma^N| = K^N.$$

To define a transition function $\Delta$, one must associate a unique next state in $\Sigma$ with each possible neighborhood state. Since there are $K = |\Sigma|$ choices of state to assign as the next state for each of the $|\Sigma^N|$ possible neighborhood states, there are $K^{(K^N)}$ possible transition functions $\Delta$ that can be defined. We use the notation $\mathcal{D}_N^K$ to refer to the set of all possible transition functions $\Delta$ which can be defined using $N$ neighbors and $K$ states.

For example, consider a two-dimensional cellular automaton using eight-states per cell, a rectangular lattice, and the five-cell neighborhood template shown above. Here $K = 8$ and $N = 5$, so $|\Delta| = K^N = 8^5$ so there are 32,768 possible neighborhood states. For each of these, there is a choice among eight states for the next cell state under $\Delta$, so there are $|\mathcal{D}_N^K| = K^{(K^N)} = 8^{(8^5)} \approx 10^{30,000}$ possible transition functions using the five-cell neighborhood template with eight-states per cell.

## 2.1 PARAMETERIZING THE SPACE OF CA RULES

$\mathcal{D}_N^K$, the set of possible $\Delta$ functions for a CA of $K$ states and $N$ neighbors, is fixed once we have chosen the number of states per cell and the neighborhood template. However, there is no intrinsic order within $\mathcal{D}_N^K$; it is a large, undifferentiated *space* of CA rules.

Imposing a *parameterization scheme* on this undifferentiated space of CA rule allows us to define a natural ordering on the rules. The ideal ordering scheme would partition the space of CA rules in such a manner that rules from the same partition would support similar dynamics. Such an ordering on $\mathcal{D}_N^K$ would allow us to observe the way in which the dynamical behaviors of CA vary from partition to partition.

The location in this space of the partitions supporting universal computation, relative to the location of partitions supporting *other* possible dynamical behaviors, would then provide us with insights into the conditions under which we should expect a complex dynamics of information to emerge in CA.

**2.1.1 THE $\lambda$ PARAMETER** We will consider a subspace of $\mathcal{D}_N^K$, characterized by the parameter $\lambda$.[26,27,28]

The $\lambda$ parameter is defined as follows. We pick an arbitrary state $s \in \Sigma$, and call it the *quiescent* state $s_q$. Let there be $n_q$ transitions to this special quiescent state in a transition function $\Delta$. Let the remaining $K^N - n_q$ transitions in $\Delta$ be filled by picking randomly and uniformly over the other $K - 1$ states in $\Sigma - s_q$. Then

$$\lambda = \frac{K^N - n_q}{K^N} \,. \tag{1}$$

If $n_q = K^N$, then *all* of the transitions in the $\Delta$ function will be to the quiescent state and $\lambda = 0.0$. If $n_q = 0$, then there will be *no* transitions to $s_q$ and $\lambda = 1.0$. When all states are represented equally in the $\Delta$ function, then $\lambda = 1.0 - 1/K$.

The parameter values $\lambda = 0$ and $\lambda = 1 - 1/K$ represent the most homogeneous and the most heterogeneous $\Delta$ functions, respectively. The behavior in which we will be interested is captured between these two parameter values. Therefore, we experiment primarily with $\lambda$ in the range $0 \leq \lambda \leq 1 - 1/K$.



$\lambda$ discriminates well between dynamical regimes for "large" values of $K$ and $N$, whereas $\lambda$ discriminates poorly for small values of $K$ and $N$. For example, for the "elementary" one-dimensional CA with $K = 2$ and $N = 3$, $\lambda$ is only very roughly correlated with dynamical behavior, which probably explains why the relationships reported here were not observed in earlier work on classifying CA dynamics,[49,51] as these investigations were carried out using CA with minimal values of $K$ and $N$.

For these reasons, we employ CA for which $K \geq 4$ and $N \geq 5$, which results in $\Delta$ functions of size $4^5 = 1024$ or larger.

**2.1.2 SEARCHING CA SPACE WITH THE $\lambda$ PARAMETER** In the following, we use the $\lambda$ parameter as a means of sampling $\mathcal{D}_N^K$ in an ordered manner. We do this by stepping through the range $0.0 \leq \lambda \leq 1.0 - 1/K$ in discrete steps, randomly constructing $\Delta$ functions for each $\lambda$ point. Then we run CA under these randomly constructed $\Delta$ functions, collecting data on various measures of their dynamical behavior. Finally, we examine the behavior of these measures as a function of $\lambda$.

$\Delta$ functions are constructed in two ways using $\lambda$: the *random-table method* and the *table-walk-through method*. In the random-table method, $\lambda$ is interpreted as a bias on the selection of states from $\Sigma$ as we sequentially fill in the transitions that make up a $\Delta$ function. To do this, we step through the table, flipping a $\lambda$-biased coin for each neighborhood state. If the coin comes up tails (with probability $1.0 - \lambda$), we assign the state $s_q$ as the next cell state for that neighborhood state. If the coin comes up heads (with probability $\lambda$), we pick one of the $K - 1$ states in $\Sigma - s_q$ at uniform random as the next cell state.

In the table-walk-through method, we start with a $\Delta$ function consisting entirely of transitions to $s_q$, so that $\lambda = 0.0$ (but note restrictions below). New transition tables with higher $\lambda$ values are generated by randomly replacing a few of the transitions to $s_q$ in the current $\Delta$ function with transitions to other states, selected randomly from $\Sigma - s_q$. Tables with *lower* $\lambda$ values are generated by randomly replacing a few transitions that are *not* to $s_q$ in the current $\Delta$ function by transitions to $s_q$.

Thus, under the table-walk-through method, we progressively perturb "the same table," whereas under the random-table method, each table is generated anew, independently of the others.

**2.1.3 FURTHER RESTRICTIONS ON CA** In order to make our studies more tractable, we impose two further conditions on the rule spaces:

1. *Strong Quiescence:* all neighborhood states uniform in cell state $s_i$ will map to state $s_i$ under $\Delta$.
2. *Spatial Isotropy:* all planar rotations of a neighborhood state will map to the same cell state under $\Delta$.

The first restriction implies that uniform regions of the array remain uniform under the action of a rule. The second restriction implies that the local physics cannot tell which way is up, so to speak. That is, local rules cannot make use of the global property of absolute orientation with respect to the lattice, although they can distinguish between left and right handedness, which is a local property.

## 2.2 CLASSIFYING CA BEHAVIOR

The most widely known scheme for classifying cellular automata on the basis of their dynamical behaviors is due to Wolfram,[49] who proposed the following four qualitative classes of CA behaviors:

- **Class I CA** evolve to a fixed, homogeneous state.
- **Class II CA** evolve to simple separated periodic structures.
- **Class III CA** yield chaotic aperiodic patterns.
- **Class IV CA** yield complex patterns of localized structures.

Wolfram finds the following analogs for his four classes of cellular automaton behaviors in the field of dynamical systems.[49]

- **Class I CA** evolve to *limit points.*
- **Class II CA** evolve to *limit cycles.*
- **Class III CA** evolve to *chaotic* behavior of the kind associated with *strange attractors.*
- **Class IV CA** have very long *transients*, and "no direct analog for them has been identified among continuous dynamical systems."

Various suites of statistical measures have been employed in the attempt to quantify these qualitative differences, with somewhat limited success.[12]

However, classification alone is not enough. What is needed is a deeper understanding of the structure of cellular automata rule spaces, one which provides an explanation for the existence of the observed classes and for their relationships to one another. Choosing an appropriate parameterization of the space of cellular automata rules, such as $\lambda$, allows direct observation of the way in which different statistical measures are related as a function of the parameter(s). These relationships in turn provide an explanation for the existence and ordering of the various qualitatively distinguishable classes of cellular automata dynamics. From the vantage point of the resulting understanding of the "deep structure" of CA rule spaces, the various qualitative classes proposed by Wolfram and others make sense, and are even to be expected. However, it is also obvious that such classification schemes can only serve as rough approximations to the more subtle, underlying structure.



### 2.2.1 CLASS IV CELLULAR AUTOMATA

Of the four Wolfram classes, Class IV is both the most interesting and the least well characterized. It is the most interesting class because it contains the CA rules exhibiting the most "complex" behaviors. Conway's game of LIFE[16] is a well-known example of a Class IV CA. Class I CA exhibit the maximal possible order: a uniform, homogeneous structure, like a crystal. This order exists on both *local* and *global* scales. Class II CA also exhibit both local and global order, although not maximal. Class III CA can exhibit maximal *disorder*, and this disorder exists on both local and global scales. Class IV CA exhibit a great deal of local order, but little apparent global order, although with time, global order may emerge.

Wolfram suggests that Class IV CA are capable of supporting computation, even universal computation, and that it is this capacity that makes their behavior so complex. The game of LIFE has been shown to be capable of universal computation.[6] In this paper, we provide support for this hypothesis and offer an explanation for why we should expect to see a capacity for information processing in certain regions of CA rule space. The association of Class IV CA with "very long transients" will figure "critically" in the explanation, and helps to explain why Class IV is hard to characterize.

# 3. QUALITATIVE OVERVIEW OF CA DYNAMICS

In this section, we present a *qualitative* overview of the change in dynamical behavior of a typical one-dimensional CA as we pass through the most interesting region of CA rule space by varying $\lambda$ via the table-walk-through method.[2] In the next section, we will present a *quantitative* overview.

## 3.1 A SURVEY OF ONE-DIMENSIONAL CA DYNAMICS USING $\lambda$

For these one-dimensional CA, $K = 4$ and $N = 5$ (i.e., *two* cells on the left and *two* cells on the right are included in the neighborhood template). The arrays consist of 128 sites connected in a circle, providing periodic boundary conditions. Each array is started from a random initial configuration on the top line, and successive lines show successive time steps in the evolution of the dynamical behavior.

For each value of $\lambda$, we show two evolutions. The arrays in Figure 1 are started from a uniform random initial configuration over all 128 sites. This series illustrates the kinds of structures that develop, as well as the typical transient times before these structures are achieved.

[2] It is not always the case that the table-walk-through method passes through the most interesting region of CA rule space. This topic will be taken up in the next section.

The arrays in Figure 2 are initialized with a patch of 20 randomized sites, with the remaining sites set to state 0 (the "quiescent" state). This series illustrates the relative rates of the spread (or collapse) of the area of dynamical activity.

For those values of $\lambda$ exhibiting long transients, we have reduced the scale of the arrays in order to display longer evolutions.

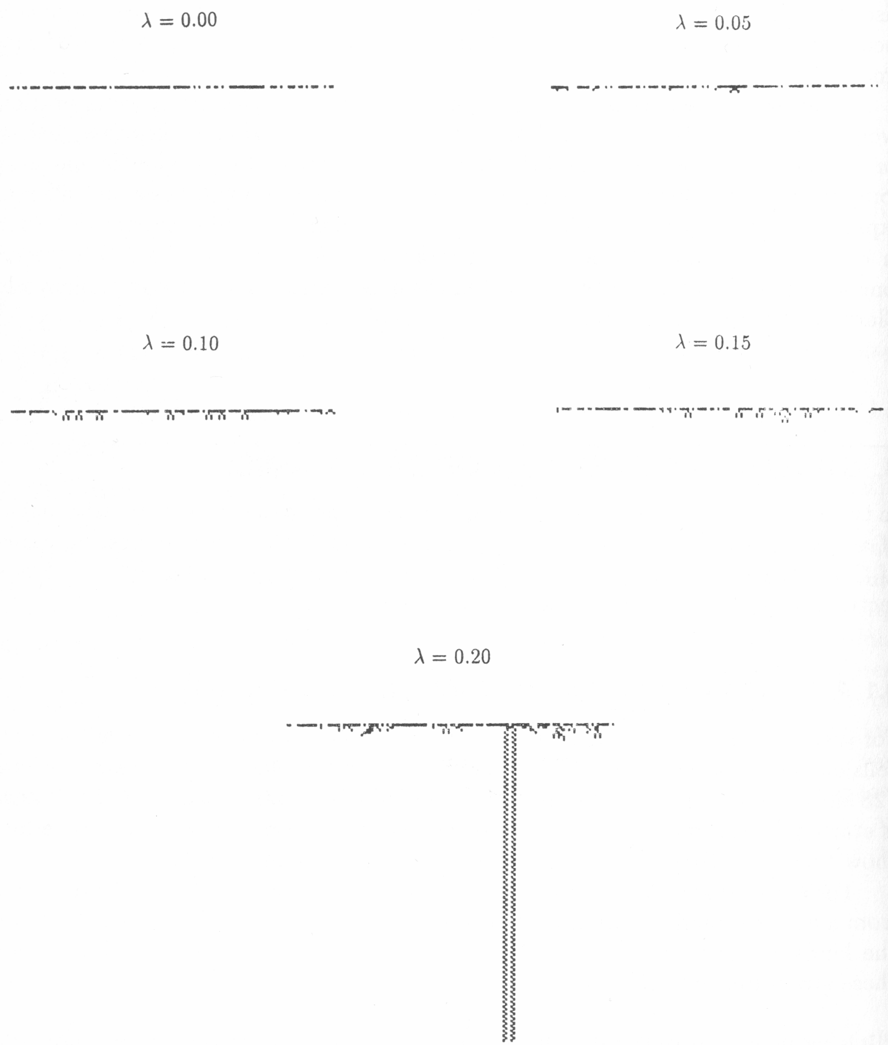


FIGURE 1 Evolution of one-dimensional CA over $0 < \lambda \leq 0.75$ from fully random initial conditions.



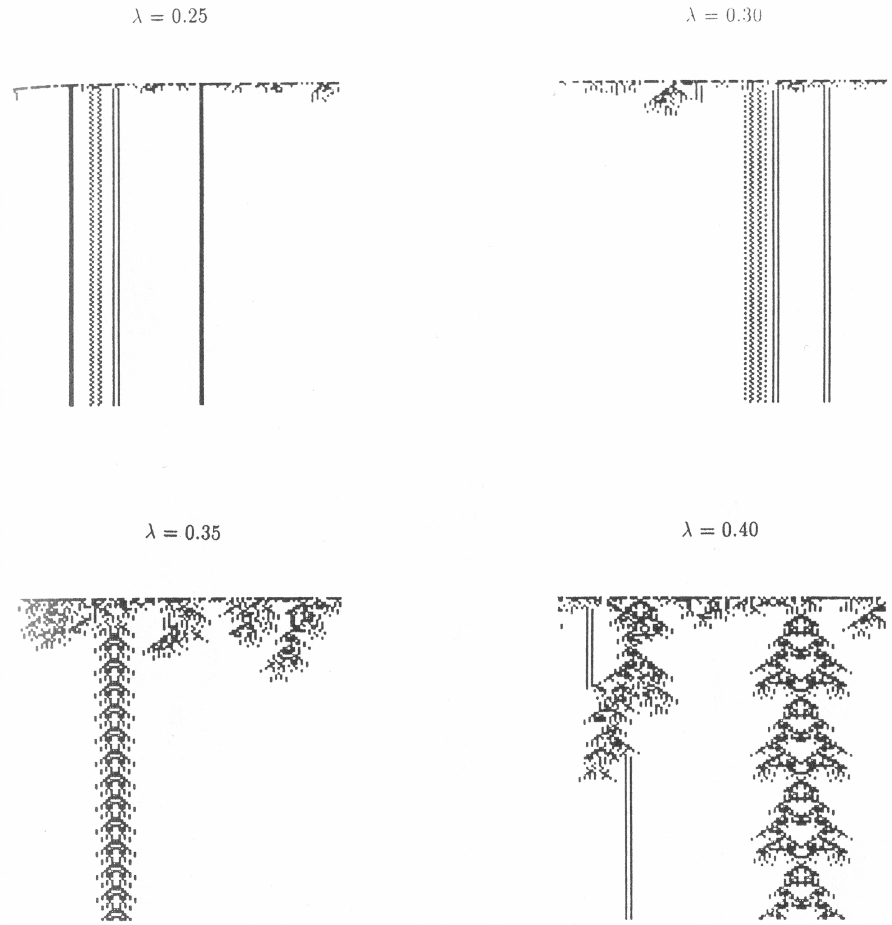


FIGURE 1 (continued)

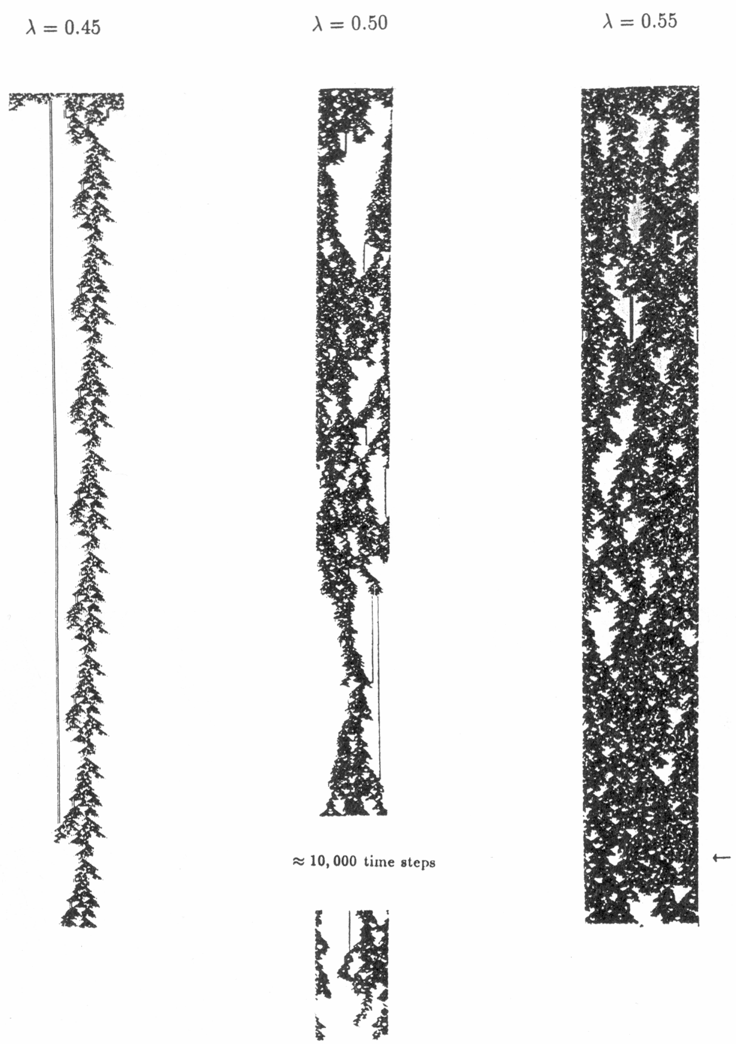


FIGURE 1 (continued)



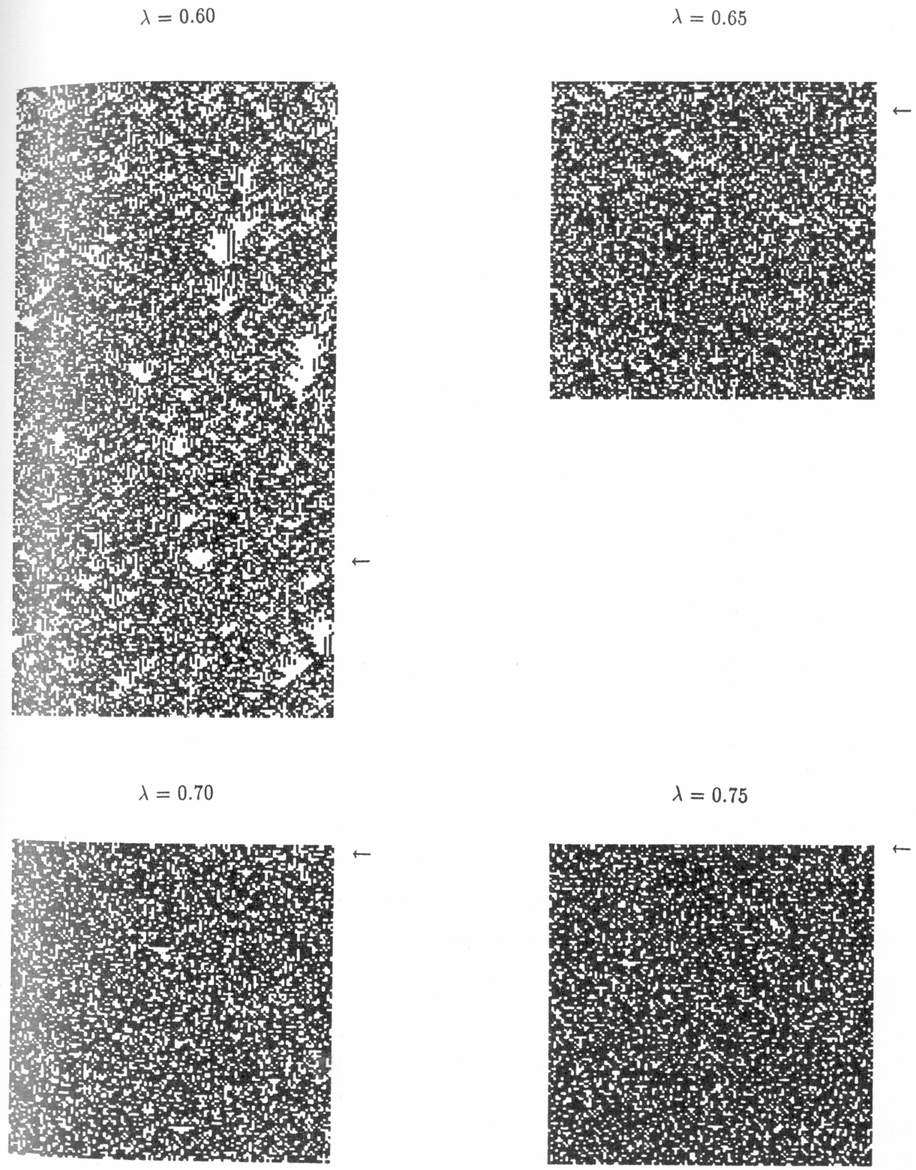


FIGURE 1 (continued)

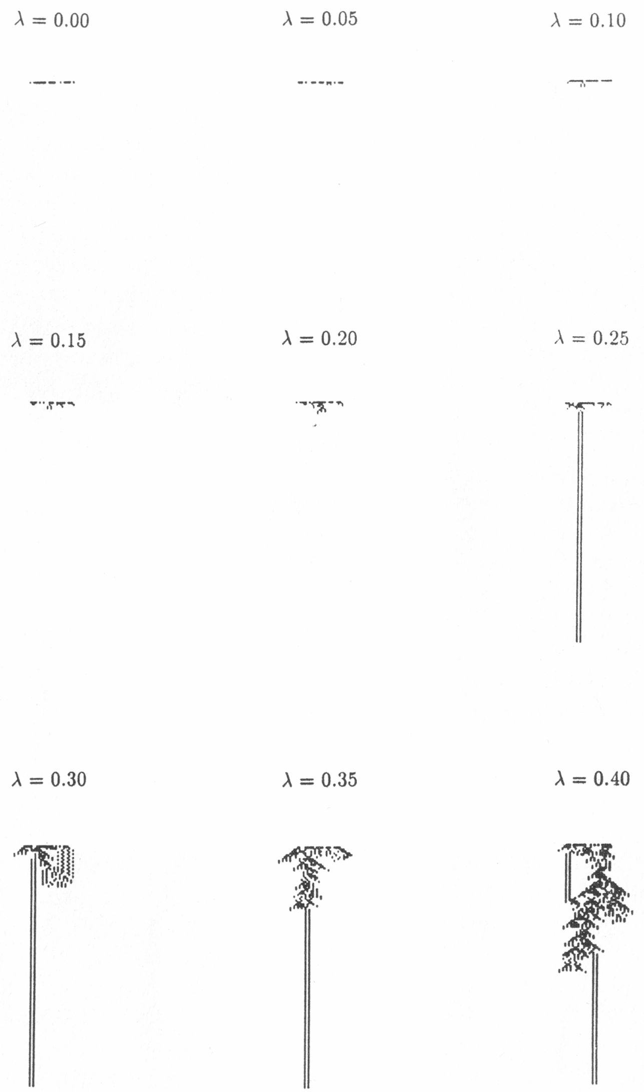


FIGURE 2 Evolution of one-dimensional CA over $0 < \lambda \leq 0.75$ from partially random initial conditions.



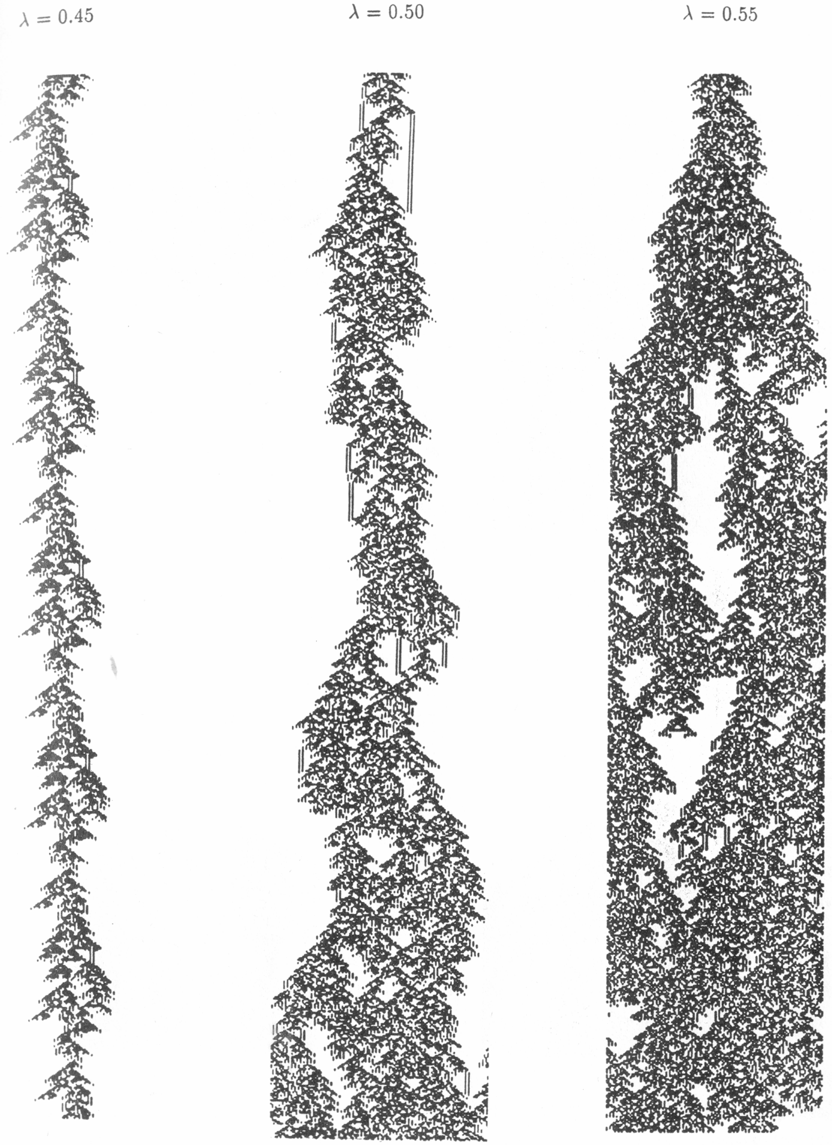


FIGURE 2 (continued)

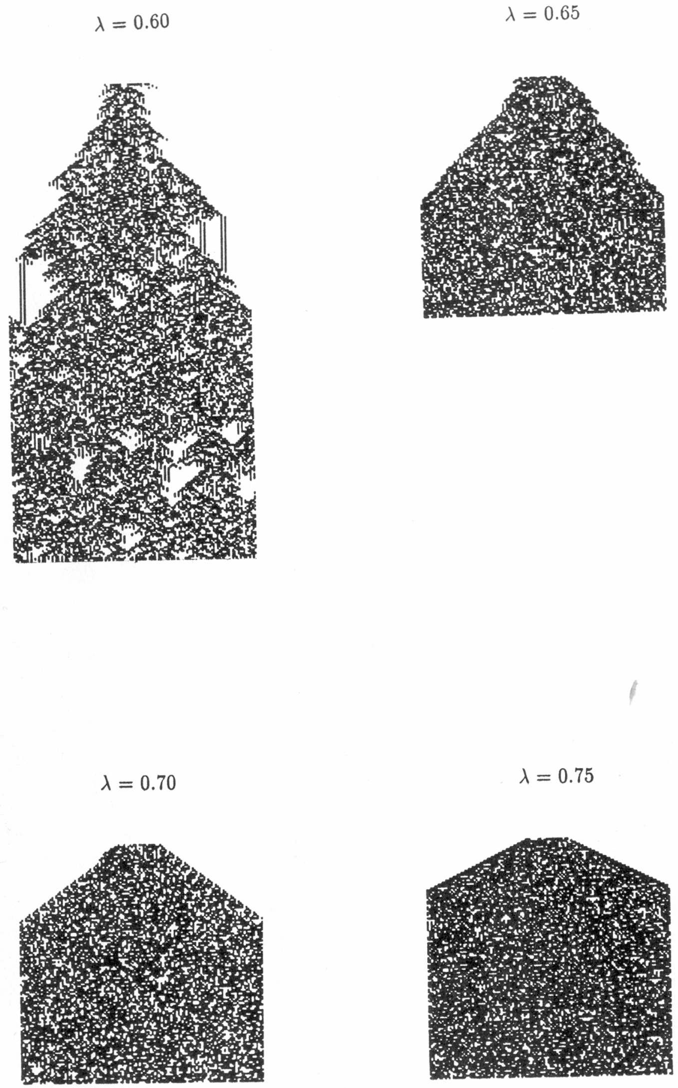


FIGURE 2 (continued)

We start with $\lambda \approx 0$. Note that under our strong quiescence condition, we cannot have $\lambda$ be identical to 0. The salient features observed as we vary $\lambda$ away from 0 are itemized below. Note that these examples are the result of a *single* table walk-through, and therefore the association of specific dynamical behaviors with specific values of $\lambda$ is not particularly significant.



$\lambda \approx 0.00$: All dynamical activity dies out after a single time step, leaving the arrays uniform in state $s_q$. The area of dynamical activity has collapsed to nothing.

$\lambda \approx 0.05$: The dynamics reaches the uniform $s_q$ fixed-point after approximately 2 time steps.

$\lambda \approx 0.10$: The homogeneous fixed point is reached after 3 or 4 time steps.

$\lambda \approx 0.15$: The homogeneous fixed point is reached after 4 or 5 time steps.

$\lambda \approx 0.20$: The dynamics reaches a periodic structure, which will persist forever (Figure 1.20). Transients have increased to 7–10 time steps as well. Note that the evolution does not necessarily lead to periodic dynamics (Figure 2.20).

$\lambda \approx 0.25$: Structures of period 1 appear. There are now three different possible outcomes for the ultimate dynamics of the system, depending on the initial state. The dynamics may reach a *homogeneous* fixed point (consisting entirely of state $s_q$,) or it may reach a *heterogeneous* fixed point (consisting mostly of cells in state $s_q$, with a sprinkling of cells stuck in the other states,) or it may settle down to periodic behavior. Notice that the transients have lengthened even more.

$\lambda \approx 0.30$: Transients have lengthened again.

$\lambda \approx 0.35$: Transient length has grown significantly, and a new kind of periodic structure with a longer period has appeared (Figure 1.35). Most of the previous structures are still possible; hence, the spectrum of dynamical possibilities is broadening.

$\lambda \approx 0.40$: Transient length has increased to about 60 time steps, and a structure has appeared with a period of about 40 time steps. Areas of dynamical activity are still collapsing down onto isolated periodic configurations.

$\lambda \approx 0.45$: Transient length has increased to almost 1,000 time steps. (Figure 1.45). Here, the structure on the right appears to be periodic, with a period of about 100 time steps. However, after viewing several cycles of its period, it is apparent that the whole structure is moving to the left, and so this pattern will not recur precisely in the same position until it has cycled at least once around the array. Furthermore, as it propagates to the left, this structure eventually annihilates a period one structure after about 800 time steps. Thus, the transient time before a cycle is reached has grown enormously. It turns out that even after one orbit around the array, the periodic structure does not return exactly to its previous position. It must orbit the array 3 times before it repeats itself exactly. As it has shifted over only 3 sites after its quasi-period of 116 time steps, the true period of this structure is 14,848 time steps. Here, the area of dynamical activity is at a balance point between collapse and expansion, as illustrated in Figure 2.45.

$\lambda \approx 0.50$: Typical transient length is on the order of 12,000 time steps. After the transient, the dynamical activity settles down to periodic behavior, possibly of period one as shown in the figure. Although the dynamics eventually becomes simple, the transient time has increased dramatically. Note in Figure 2.50 that the general tendency now is that the area of dynamical activity *expands* rather than contracts with time. There are, however, large fluctuations in the area covered by dynamical activity, and it is these fluctuations that lead to the eventual collapse to periodic dynamics in a finite array.

$\lambda \approx 0.55$: Whereas before, the dynamics *eventually* settled down to periodic behavior, we are now in a regime in which the dynamics settles down to effectively *chaotic* behavior. Furthermore, the previous trend of transient length *increasing* with increasing $\lambda$ is reversed. The arrow to the right of the evolutions in Figures 1.55–1.75 indicates the approximate time by which the average site occupation density has settled down to within 1% of its long-time average. Note that the area of dynamical activity expands more rapidly with time.

$\lambda \approx 0.60$: The dynamics are quite chaotic, and the transient length to "typical" chaotic behavior has decreased significantly. The area of dynamical activity expands more rapidly with time.

$\lambda \approx 0.65$: Typical chaotic behavior is achieved in only ten time steps or so. The area of dynamical activity is expanding at about one cell per time step in each direction, approximately half of the maximum possible rate for this neighborhood template.

$\lambda \approx 0.70$: Fully developed chaotic behavior is reached in only two time steps. The area of dynamical activity is expanding even more rapidly.

$\lambda \approx 0.75$: After only a single time step, the array is essentially random and remains so thereafter. The area of dynamical activity spreads *at* the maximum possible rate.

Therefore, by varying the $\lambda$ parameter throughout $0.0 < \lambda \leq 0.75$ over the space of possible $K = 4$, $N = 5$, one-dimensional cellular automata, we progress from CA exhibiting the maximal possible order to CA exhibiting the maximal possible disorder. At intermediate values of $\lambda$, we encounter a *phase transition* between periodic and chaotic dynamics, and while the behavior at either end of the $\lambda$ spectrum seems "simple" and easily predictable, the behavior in the vicinity of this phase transition seems "complex" and unpredictable.

## 3.2 COMMENTS ON QUALITATIVE DYNAMICS

A few comments on these examples are in order.

First, the progression through the spectrum of dynamical behaviors as a function of $\lambda$ is clearly:

$$\text{fixed-point} \rightarrow \text{periodic} \rightarrow \text{"complex"} \rightarrow \text{chaotic}.$$

In terms of the Wolfram classes, the sequence is:

$$I \rightarrow II \rightarrow IV \rightarrow III. \tag{2}$$

That is, the complex rules are located *inbetween* the periodic and the chaotic rules.

Second, there is clearly a phase transition between periodic and chaotic behavior, with $\lambda_c \approx 0.50$ (in these examples).



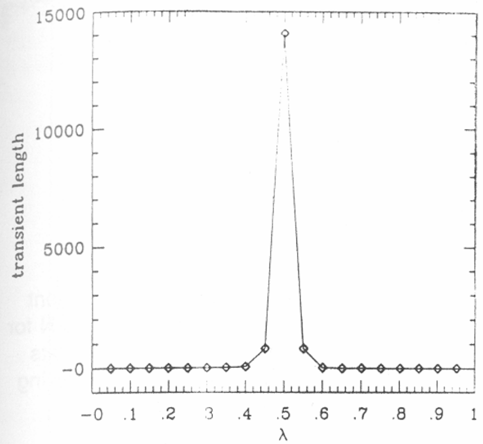


FIGURE 3 Average transient length T versus $\lambda$. Transient length apparently diverges rapidly in the vicinity of the transition.

Third, transients are short at either end of the range of $\lambda$ and get longer as we approach the middle of the range.[3]

Transients clearly diverge in the vicinity of $\lambda_c$, evidence of the second-order phenomenon of *critical slowing down*. As we continue to raise $\lambda$ beyond $\lambda_c$, although the dynamics are now settling down to effectively chaotic behavior instead of periodic behavior, the transient lengths are getting *shorter* with increasing $\lambda$, rather than longer. The relationship between transient length and $\lambda$ for these examples is plotted in Figure 3.

Fourth, clearly, in CA for which $\lambda \ll \lambda_c$ or $\lambda \gg \lambda_c$, we do not have to wait very long before we can be almost certain about the ultimate outcome of the dynamics. For CA with $\lambda \approx \lambda_c$, however, the dynamics will look very much the same, whether it will ultimately result in a periodic state, or will remain non-periodic "forever." In general, the outcome is undecidable. We will discuss this further in relation to Turing's famous Halting problem.

Fifth, the size of the array has an effect on the dynamics only for intermediate values of $\lambda$. For low values of $\lambda$, array size has *no* discernible effect on transient length. Not until, $\lambda \approx 0.45$, do we begin to see a small difference in the transient length as the size of the array is increased. For $\lambda = 0.50$, however, array size has a significant effect on transient length. The growth of transient length as a function of array size for $\lambda = 0.50$ is plotted in Figure 4. The essentially linear relationship on this log-normal plot suggests that transient length depends *exponentially* on array size for $\lambda \approx 0.50$.

[3]Note that "transients" for chaotic CA are defined differently than for periodic CA. Transients for periodic CA are defined in terms of the number of iterations before a cycle is reached, while transients for chaotic CA are defined in terms of iterations before statistical convergence is achieved. This differs from the traditional CA literature, but is more in line with the literature on dynamical systems and chaos theory.

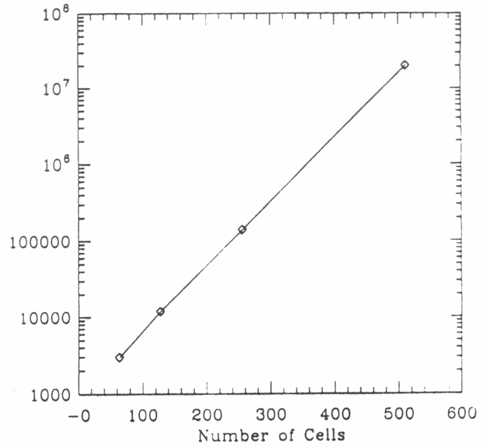


FIGURE 4 Average transient length T versus array size N for $\lambda = 0.50$. This plot suggests that transient length is growing exponentially with N in the vicinity of $\lambda_c$.

Furthermore, transient times exhibit *decreasing* dependence on array size as $\lambda$ is increased past the transition point. For the highest value of $\lambda$ (= 0.75 in this case) when all states are represented uniformly in the transition table, the transient lengths exhibit *no* dependence on array size, as was the case for the lowest values of $\lambda$. We will discuss this further when we use this fact to help explain the existence of computational complexity classes.

Sixth, the behavior of the CA dynamics is most complicated in the vicinity of the transition. Compare the length of the description of the dynamics at each $\lambda$ value in the previous section. It takes longer to describe what is going on near the transition than it does to describe what is going on far from the transition. The dynamics becomes simply periodic for low $\lambda$, whereas for high $\lambda$ the randomness simply spreads outwards in a uniformly expanding "circle" at the maximum possible rate. The mutual information and entropy data presented in the next section will quantify the important distinction between complexity and randomness.

Finally, it is important to note that the transition region supports both static and propagating structures (Figure 1.45). The propagating structures are essentially *solitary waves*, quasi-periodic patterns of state change, which—like the "glider" in Conway's game of LIFE[16]—propagate through the array, constantly moving with respect to the fixed background of the lattice. Figure 5 traces the time evolution of an array of 512 sites, and shows that the rule governing the behavior of Figure 1.45 supports several different kinds of "particles," which interact with each other and with the static periodic structures in complicated ways (e.g., note the manner in which the collision of a propagating particle with a static periodic structure produces a particle traveling in the opposite direction). Such propagating and static structures can form the basis for signals and storage, and interactions between them can modify either stored or transmitted information in the support of an overall computational process.



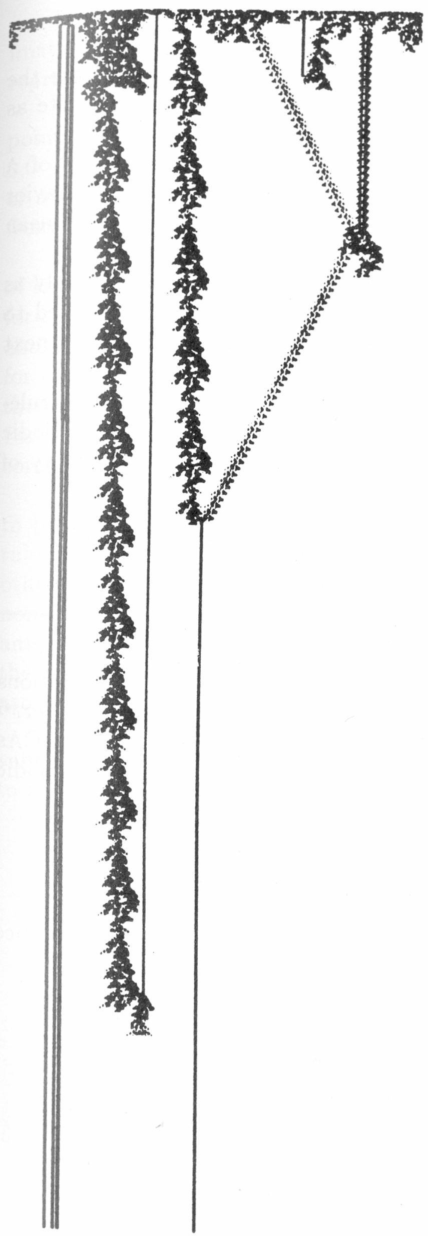

FIGURE 5 Evolution of a one-dimensional array consisting of 512 sites under the rule governing the evolution illustrated in Figure 1, $\lambda = .45$.

### 3.3 QUALIFICATIONS

It must be pointed out that although the examples presented here illustrate the most interesting change in dynamics as a function of $\lambda$, the story is not quite as simple as we have presented it here.

The situation is complicated by two factors. First, different traversals of $\lambda$ space using the table-walk-through method make the transition to chaotic behavior at different $\lambda$ values, although there is a well-defined distribution around a mean value.

Second, one does not always pass through the "complex" regime as neatly as in this example. More often, the dynamics jumps directly from fairly ordered to fairly disordered behavior. We will discuss both of these complications in the next section.

Despite these qualifications, the overall picture is clear: as we survey CA rule-spaces using the $\lambda$ parameter, we encounter a phase transition between periodic and chaotic behavior, and the most complex behavior is found in the vicinity of this transition.

## 4. QUANTITATIVE OVERVIEW OF CA DYNAMICS

In this section, we present a brief quantitative overview of the structural relations among the dynamical regimes in CA rule spaces as revealed by the $\lambda$ parameter.[4]

The results of this section are based on experiments using two-dimensional CAs with $K = 8$ and $N = 5$. Arrays are typically of size 64 x 64, and again, periodic boundary conditions are employed.

### 4.1 MEASURES OF COMPLEXITY

The measures employed were chosen for their collective ability to reveal the presence of information in its various forms within CA dynamics.

4.1.1 SHANNON ENTROPY We use Shannon's Entropy $H$ to measure basic information capacity. For a discrete process $A$ of $K$ states[5]:

$$H(A) = -\sum_{i=1}^{K} p_i \log p_i \,. \tag{3}$$

[4]For a more detailed review, see Langton.[28]

[5]Throughout, log is taken to the base 2; thus, the units are bits.



Figure 6 shows the average entropy per cell, $\overline{H}$, as a function of $\lambda$ for approximately 10,000 CA runs. The random-table method was employed, so each point represents a distinct random transition table.

First, note the overall envelope of the data and the large variance at most $\lambda$ points. Second, note the sparsely populated gap over $0.0 \le \lambda \lesssim 0.6$ and between $0.0 \le \overline{H} \lesssim 0.84$. This distribution appears to be bimodal, suggesting the presence of a phase transition. Third, note the rapid decrease in variability as $\lambda$ is raised from $\sim 0.6$ to its maximum value of 0.875.

Two other features of this plot deserve special mention. First, the abrupt cutoff of low $\overline{H}$ values at $\lambda \approx 0.6$ corresponds to the *site-percolation* threshold $P_c \approx 0.59$ for this neighborhood template. Thus, we may suppose that, since $\lambda$ is a dynamical analog of the site occupation probability $P$, the *dynamical* percolation threshold for a particular neighborhood template is bounded above by the *static* percolation threshold $P_c$. This is borne out by experiments with other neighborhood templates. For instance, the nine-neighbor template exhibits a sharp cutoff at $\lambda \approx 0.4$, which corresponds well with the site percolation threshold $P_c \approx 0.402$ for this lattice.

The second feature is the "ceiling" of the gap at $\overline{H} \approx 0.84$. This turns out to be the average entropy value for one of the most commonly occurring chaotic rules. In such rules the dynamics has collapsed onto only two states—$s_q$ and one other—and the rule is such that a mostly quiescent neighborhood containing one non-quiescent state maps to that non-quiescent state. In one-dimensional cellular automata, such rules gives rise to the familiar triangular fractal pattern known as the Serpinski Gasket. There are many ways to achieve such rules, and they can be achieved at very low $\lambda$ values. Most of the low-$\lambda$ chaotic rules are of this type.

The entropy data of Figure 6 suggest an anomaly at intermediate parameter values, possibly a phase transition between two kinds of dynamics. Since there seems to be a discrete jump between low and high entropy values, the evidence points to

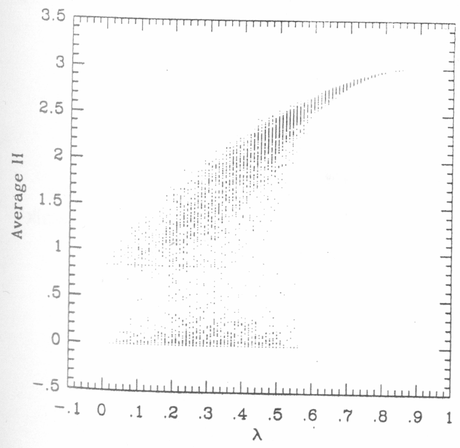


FIGURE 6 Average single cell entropy $\overline{H}$ over $\lambda$ space.

a first-order transition, similar to that observed between the solid and fluid phases of matter. However, the fact that the gap is not completely empty suggests the possibility of second-order transitions as well.

The table-walk-through method of varying $\lambda$ reveals more details of the structure of the entropy data. Figure 7 shows four superimposed examples of the change in the average cell entropy as we vary the $\lambda$ value of a table. Notice that in each of the four cases, the entropy remains fairly close to zero until—at some critical $\lambda$ value—the entropy jumps to a higher value, and proceeds fairly smoothly towards its maximum possible value as $\lambda$ is increased further. Such a discontinuity is a classic signature of a first-order phase transition. Most of our complexity measures exhibit similar discontinuities at the same $\lambda$ value *within a particular table.*

Notice also that the $\lambda$ value at which the transition occurs is different for each of the four examples. Obviously, the same thing—a jump—is happening as we vary $\lambda$ in each of these examples, but it happens at different values of $\lambda$. When we superimpose 50 runs, as in Figure 8, we see the internal structure of the entropy data envelope plotted in Figure 6.

Since we have located the transition events, we may line up these plots by the events themselves, rather than by $\lambda$, in order to get a clearer picture of what is going on before, during, and after the transition. This is illustrated in Figure 9. The abcissa is now measured in terms of $\Delta\lambda$: the distance from the transition event. Figure 10 shows the same data as Figure 8 but lined up by $\Delta\lambda$.

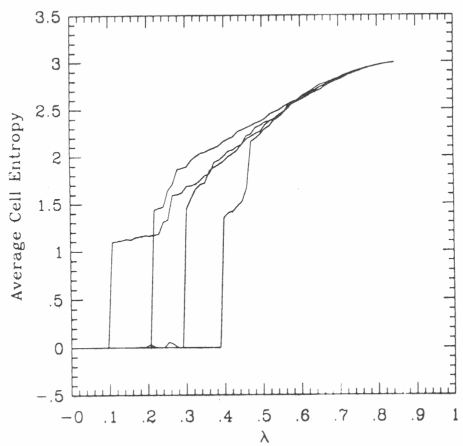


FIGURE 7 Superposition of entropy versus $\lambda$ curves for four separate runs using the table-walk-through method.



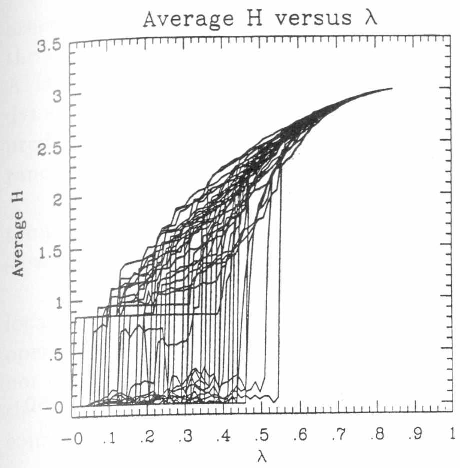


FIGURE 8 Superposition of entropy versus $\lambda$ curves for 50 separate runs using the table-walk-through method.

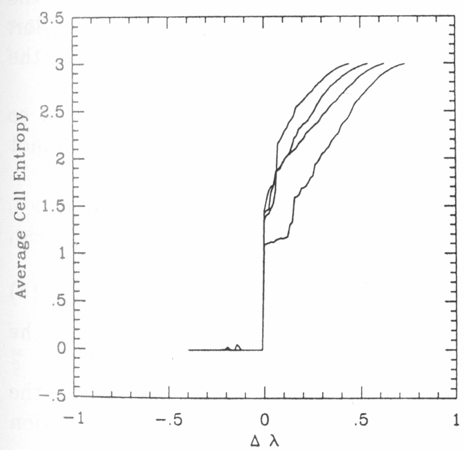


FIGURE 9 Superposition of entropy versus $\lambda$ curves for the four runs plotted in Figure 7, lined up by $\Delta\lambda$.

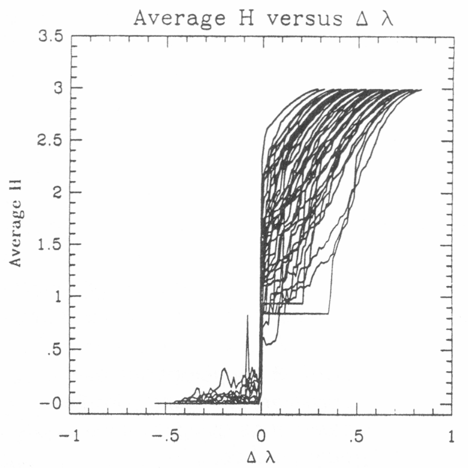


FIGURE 10 Superposition of entropy versus $\lambda$ curves for the 50 runs plotted in Figure 8, but lined up by $\Delta\lambda$.

**4.1.2 MUTUAL INFORMATION** In order for two distinct cells to cooperate in the support of a computation, they must be able to affect one another's behavior. Therefore, we should be able to find correlations between events taking place at the two cells.

The mutual information $I(A;B)$ between two cells A and B can be used to study correlations in systems when the values at the sites to be measured cannot be ordered, as is the case for the states of the cells in cellular automata.[30]

The mutual information is a simple function of the individual cell entropies, $H(A)$ and $H(B)$, and the entropy of the two cells considered as a joint process, $H(A,B)$:

$$I(A;B) = H(A) + H(B) - H(A,B) \tag{4}$$

This is a measure of the degree to which the state of cell A is correlated with the state of cell B, and *vice versa*.

Figure 11 shows the average mutual information between a cell and itself at the next time step. Note the tight convergence to low values of the mutual information for high $\lambda$ and the location of the highest values.

The increase of the mutual information in a particular region is evidence that the correlation length is growing in that region, further evidence for a phase transition.

Figure 12 shows the behavior of the average mutual information as $\lambda$ is varied, both against $\lambda$ and $\Delta\lambda$. The average mutual information is essentially zero below the transition point, it jumps to a moderate value at the transition, and then decays slowly with increasing $\lambda$. The jump in the mutual information clearly indicates the onset of the chaotic regime, and the decaying tail indicates the approach to effectively random dynamics. The lack of correlation between even adjacent cells at high $\lambda$ values means that cells are *acting* as if they were independent of each other, even though they are causally connected. The resulting global dynamics is the same as if each cell picked its next state at uniform random from among the $K$ states, with no consideration of the states of its neighbors. This kind of global dynamics is predictable in the same statistical sense that an ideal gas is globally predictable. In fact, it is appropriate to view this dynamical regime as a hot gas of randomly flipping cells.



Figure 13 shows the average mutual information curves for several different temporal and spatial separations. Note that the decay in both time and space is slowest in the middle region.

At intermediate $\lambda$ values, the dynamics support the preservation of information locally, as indicated in the peak in correlations between distinct cells. If cells are cooperatively engaged in the support of a computation, they must exhibit some—but not *too* much— correlation in their behaviors. If the correlations are too strong, then the cells are overly dependent, with one mimicking the other—not a cooperative computational enterprise.

On the other hand, if the correlations are too small, then the cells are overly *independent*, and again, they cannot cooperate in a computational enterprise, as each cell does something totally unpredictable in response to the state of the other. Correlations in behavior imply a kind of common code, or protocol, by which changes of state in one cell can be recognized and understood by the other as a *meaningful signal*. With no correlations in behavior, there can be no common code with which to communicate information.

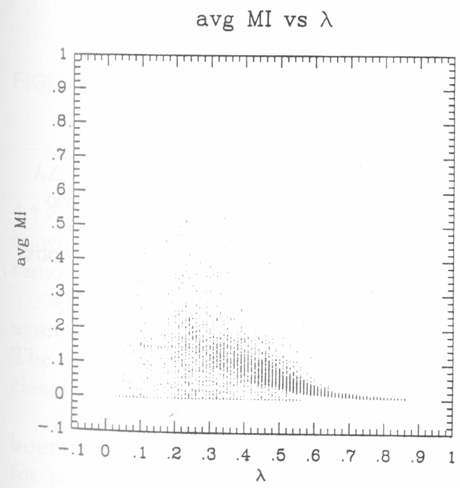


FIGURE 11 Raw data on mutual information between a cell and itself at the next time step plotted over $\lambda$.

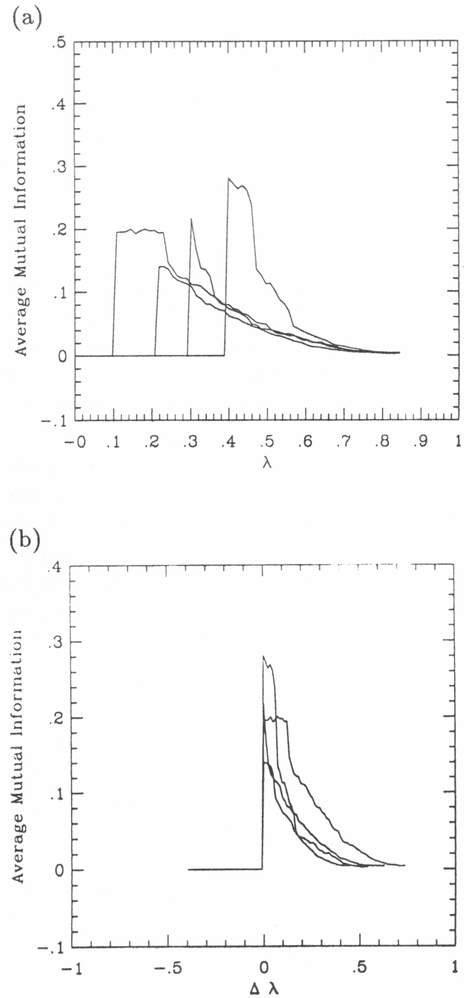


FIGURE 12 Plots of mutual information versus a) $\lambda$ and b) $\Delta\lambda$ for four runs using the table-walk-through method. The steep rise indicates the transition from periodic to chaotic behavior.



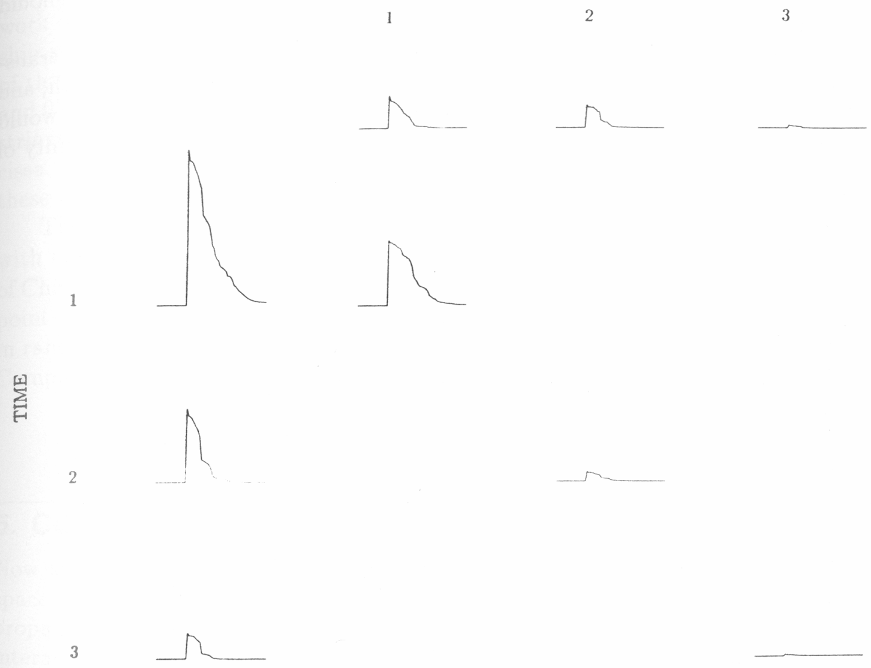


FIGURE 13 Decay of mutual information in space and time.

4.1.3 MUTUAL INFORMATION AND ENTROPY It is often useful to examine the way in which observed measures behave when plotted against one another, effectively removing the (possibly unnatural) ordering imposed by the control parameter.

Of the measures we have looked at, the most informative pair when plotted against each other are the mutual information and the average single cell entropy. The relationship between these two measures is plotted in Figure 14. Again, we see clear evidence of a phase transition.

The envelope of the relationship is bounded below the transition by the linear bound that $H$ places on the mutual information. All of the points on this line are for periodic CAs. This line intersects the curve bounding the envelope *above* the transition at an entropy value $H_c \approx 0.32$ on the normalized entropy scale.

This is a *very* informative plot. There is a clear, sharply defined maximum value of mutual information at a specific value of the entropy, and the mutual information falls off rapidly on either side. This seems to imply that there is an *optimal working* *entropy* at which CAs exhibit large spatial and temporal correlations. Why should this be the case?

Briefly, information storage involves *lowering* entropy while information transmission involves *raising entropy*.[18] In order to compute, a system must do both, and therefore must effect a tradeoff between high and low operating entropy. It would seem from the work reported here that this tradeoff is optimized in the vicinity of a phase transition.

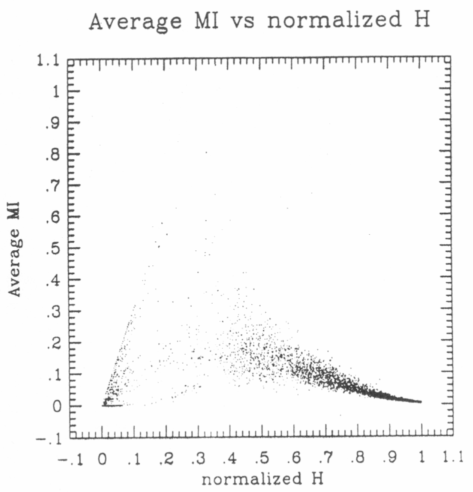


FIGURE 14 Mutual information versus normalized entropy.

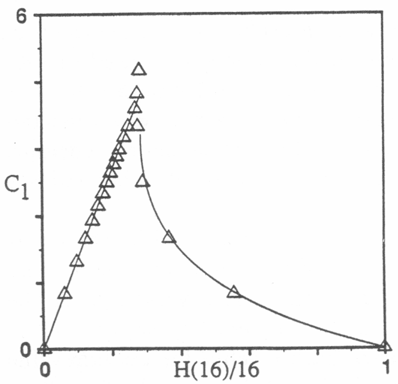


FIGURE 15 Crutchfield's plot of complexity versus normalized entropy for the logistic map (from Crutchfield.[11])



A similar relationship has been observed by Jim Crutchfield at Berkeley in his work on the transition to chaos in continuous dynamical systems.[11] This relationship is illustrated in Figure 15. Briefly, the ordinate of this plot—$C$—is a measure of the size of the minimal finite state machine required to recognize strings of 1's and 0's generated by a dynamical system (the logistic map, in this case) when these strings are characterized by the normalized per-symbol entropy listed on the abscissa. The observance of this same fundamental entropy/complexity relationship in these different classes of dynamical systems is very exciting.

These relationships support the view that, rather than increasing monotonically with randomness—as is the case for the usual measures of complexity, such as that of Chaitin and Kolmogorov[9,25]—complexity increases with randomness only up to a point—*a phase transition*—after which complexity *decreases* with further increases in randomness, so that total disorder is just as "simple," in a sense, as total order. Complex behavior involves a mix of order and disorder.

## 5. COMPLEX CA AND COMPUTATION

Now that we have demonstrated the existence of a "complex" domain within the space of CA rules—a domain that typically supports complex interactions between propagating and static structures—what reasons do we have to believe that such interactions could support complex information processing, even universal computation?

This is the question of *constructability:* Can these complex interactions be pressed into useful service as logical building blocks in the construction of a universal computing device? In this section, we provide evidence that the dynamics typically exhibited by complex CA rules can be applied to the construction of embedded general-purpose computers.

### 5.1 PROVING UNIVERSALITY OF CA RULES

There can be no universal algorithm for deciding whether an arbitrary CA rule will support universal computation. The only way to prove such a capacity for a particular rule is by construction. If one can construct a general-purpose computer under a particular CA rule—or show that all of the necessary components can be constructed and that they could be wired together in principle—then that is a sufficient proof. Failure to find such a construction does *not* constitute a proof that a particular rule will *not* support universal computation; it merely leaves the question open.

The thing that makes this problem hard is that the dynamics of interactions between simple elements supported by complex rules are often more reminiscent of chemistry than of logic. Figuring out how to go about constructing logical elements from an arbitrary given rule is like figuring out how to build logical elements out of interactions between nucleotides and proteins. It is clearly possible that one could build a biological universal computer,[6] however, designing computers based on an arbitrary "chemical" logic is much more difficult than designing computers based on more familiar logical elements.

In short, logical functions such as **AND**, **OR**, and **NOT** are typically not primitives under an arbitrary rule, and it may be a very difficult task to find how to construct them out of the functions that ARE primitive under a particular rule.

Constructive proofs of this type often have two stages. In the first stage, one attempts to incorporate the set of low-level interactions supported by a complex rule into a basic set of more familiar logical switches and elements, such as **AND**, **OR**, and **NOT** gates, clocks, wire-crossings, and etc. In the second stage, one works with these more familiar logical switches to construct an embedded computer, which might be a general-purpose RAM computer, a universal Turing Machine, or an instance of some other class of universal computational devices.

This is not to say that computation could not be taking place under a particular rule unless its dynamics support **AND**, **OR**, and **NOT** gates. Computation might be taking place based on a very different set of logical primitives naturally supported by a particular rule. It might be very difficult to figure out exactly what those primitives are. However, most constructive proofs first show that things like **AND**, **OR**, and **NOT** gates can be constructed because those are the logical primitives that we are familiar with and know how to assemble into universal computers. Thus, such proofs simply demonstrate a *capacity* for universal computation, which might be realized much more "naturally" in a very different manner under a particular rule, but which we might not recognize as constituting universal computation.

In the following sections, we will demonstrate the "constructability" of a complex rule by showing how some of the logical switches involved in the first stage discussed above may be built. A complete construction involving the second stage is beyond the scope of this paper. Here, I merely want to provide convincing evidence that such a construction could be carried through.

#### 5.1.1 CATALOGUING THE BASIC INTERACTIONS

The first stage involves performing a large set of experiments to determine all of the most basic interactions between the simple structures supported by a rule.

Sometimes these interactions have simple results (e.g., a "glider" is reflected or a "blinker" is displaced). At other times, these interactions may have more complicated results (such as the production of more gliders or blinkers), and the site of the interactions may exhibit complex activity for some time, generating lots of by-products in the process.

Once these interactions and their products have been carefully catalogued, the experimenter has a set of reactions to work with, and can try to arrange configurations in a CA so that the products of one reaction become the inputs to other reactions, on just the right trajectories, and with just the right timing. In other

[6] A human brain with the ability to mark and read the environment constitutes an existence proof.



words, one has to work with the "material" at hand, and try to mold its natural dynamics and properties into the kind of computational building blocks that are to be used in the construction of a general-purpose computational device.

### 5.2 THE PROOF THAT LIFE IS UNIVERSAL

The game of LIFE is a complex rule that has been proven capable of supporting universal computation via a constructive proof of the kind discussed above.[6] $\lambda_{LIFE} \approx 0.273$, which locates it near the critical regime for $K = 2$, $N = 9$ CA rules.

In this section, we will review this constructive proof. In the next section, we will show that the same kind of construction can be carried out for rules in the complex regime, using one of the complex rules generated by the table-walk-through method as an example.

In the proof of the universality of the game of LIFE, a stream of regularly spaced, propagating structures called "gliders" represents a string of bits. In such a stream, the presence of a glider at a potential bit-position represents a "1" while the absence of a glider at a bit-position (a "hole") represents a "0." The **AND**, **OR**, and **NOT** gates are implemented by colliding together such streams of gliders and holes.

An important configuration in the game of LIFE is the *Glider Gun:* a periodic configuration that produces a steady stream of regularly spaced gliders, i.e., a steady

FIGURE 16 A Glider Gun, a periodic LIFE configuration which produces a steady stream of gliders.

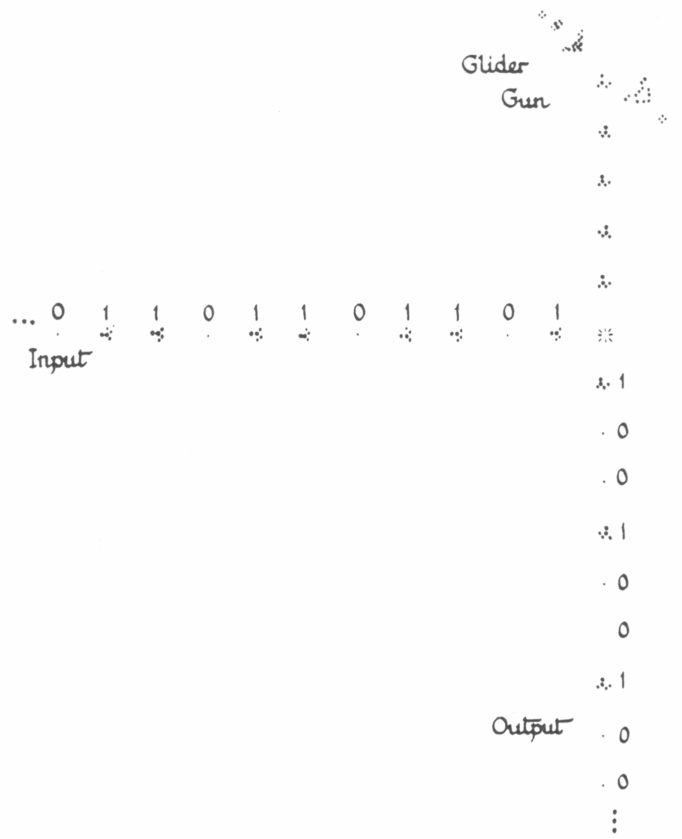


FIGURE 17 A NOT gate in the game of LIFE (from Berlekamp[6]).

stream of "1's." A glider gun producing a stream of gliders is shown in Figure 16. Such periodic "all-1" bit streams are useful for timing and for providing fixed inputs to logical elements.

Figure 17 shows the construction of a **NOT** gate in the game of LIFE. The output of a Glider Gun is collided at right angles with a stream of gliders encoding the input to the **NOT** gate. The two streams are timed so that when two gliders collide, they annihilate each other. Consequently, wherever there is a glider (a 1) in the input stream, it collides with a glider from the Glider Gun and they annihilate, leaving a hole (a 0) in the output stream. Wherever there is a hole in the input stream, the associated glider in the stream from the Glider Gun passes through untouched into the output stream. Therefore, the output stream contains a glider at every position where there was a hole in the input stream, and contains a hole at every position where there was a glider in the input stream. That is, the output



stream contains a 1 for every 0 in the input stream, and a 0 for every 1 in the input stream, thus implementing the **NOT** function.

The **AND** gate is illustrated in Figure 18. **AND** builds on the **NOT** function by using the output of **NOT** to gate a second input stream. $A \wedge B$ is implemented as follows. Input stream $B$ is fed into the **NOT** gate. What comes out is $\neg B$. If we now collide this $\neg B$ stream with the stream representing $A$, holes in the $\neg B$ stream will allow passage of gliders in the $A$ stream, and gliders in the $\neg B$ stream will annihilate gliders in the $A$ stream, leaving holes. Thus, the stream $A$ gated by stream $\neg B$ contains a glider only at positions where stream $A$ had a glider and stream $\neg B$ had a hole, which was wherever stream $B$ had a glider. That is, the output stream contains a 1 at every location where both inputs were 1, and has 0's everywhere else, thus implementing the **AND** function.

The **OR** gate is illustrated in Figure 18. **OR** combines the **AND** gate with another **NOT** gate as follows. The *other* stream coming out of the **AND** gate described above is just stream $\neg B$ gated by stream $A$. It has a hole wherever $\neg B$ had a hole, and in addition it has a hole wherever $\neg B$ had a glider *and* $A$ had a glider. It has a glider wherever stream $\neg B$ had a glider and stream $A$ had a hole. Thus, it implements the logical function $(\neg A \wedge \neg B)$, which simplifies to $\neg(A \vee B)$ by DeMorgan's law. The output from this is then inverted by a **NOT** gate, which yields $(A \vee B)$.

There are further elements required to construct a complete working computer, such as elements that will turn a signal stream by right angles, introduce delays for timing, implement an extendable memory, and so forth. We will not go into the details of these here. The interested reader can consult Berlekamp.[6]

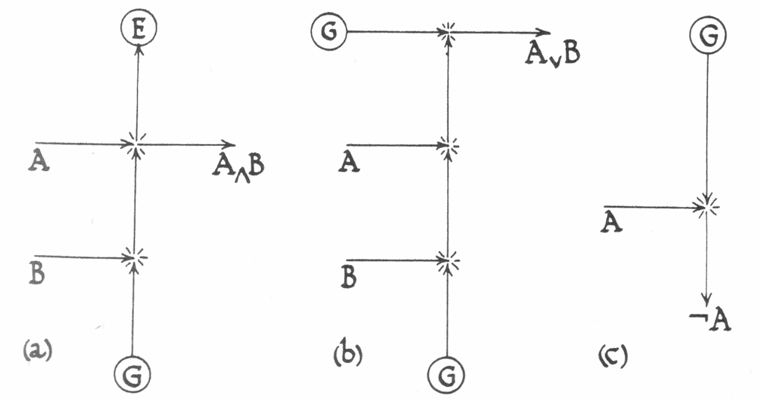


FIGURE 18 AND, OR, and NOT gates in the game of LIFE (from Berlekamp[6]).

### 5.3 CONSTRUCTABILITY FOR OTHER COMPLEX RULES

Constructions like those for **NOT**, **AND**, and **OR** above clearly demonstrate "constructability" for the game of LIFE. In this section, we show that similar constructive proofs can be provided for other complex rules by demonstrating the construction of **NOT**, **AND**, and **OR** gates within a particular $K = 2$, $N = 9$ complex rule generated via our table-walk-through method. Like the game of LIFE, this rule supports a rich variety of propagating gliders, glider guns, periodic blinkers, and a complex "chemistry" of interactions between such structures.

Figure 19 shows a Glider Gun under this rule producing a stream of gliders.[7] By colliding together streams of gliders and holes representing inputs with streams of gliders produced by glider guns and with each other, we can construct the same **AND**, **OR**, and **NOT** gates exhibited above for the game of LIFE. The **NOT** gate is shown in Figure 20. The other two gates are built up from this gate in exactly the same manner as they were in the game of LIFE.

FIGURE 19 Glider gun under a randomly generated "complex" rule.

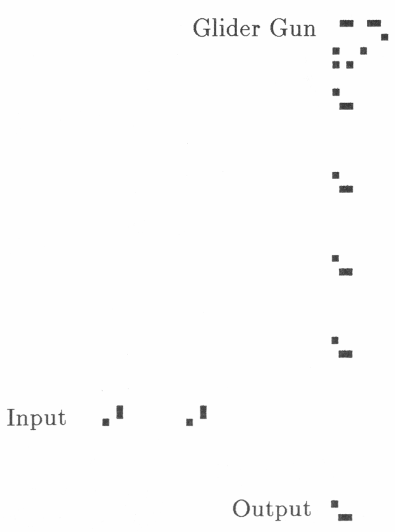


FIGURE 20 NOT gate under a "complex" rule.

[7]This particular Glider Gun propagates through space, i.e., it is a "Gliding" Glider Gun.



Again, other components will be needed to construct a complete working computer, and a proof of universality will require assembling these components together with the proper signal routing and delay elements, but the existence of these gates demonstrates "constructability" for this sample complex rule.

Similar constructions can be made for other complex rules, but not all complex rules tried have yielded to such simple constructions. In these latter cases, some aspect of the necessary componentry or interactions proved elusive. For instance, for some complex rules, none of the possible glider collisions resulted in the annihilation of both gliders. In other rules, Glider Guns were not found. These negative results do not prove that these rules are not capable of supporting universal computation—they merely suggest that this particular construction technique may not go through for these rules, whereas it is entirely possible that some other construction would go through.

## 6. PHASE TRANSITIONS AND COMPUTATION

In this section, we summarize and discuss the main points that can be derived from the evidence provided in the last several sections.

First, we review the phase-transition structure of CA rule space implied by the evidence from the previous sections, and show how this "deep structure" explains the surface-level phenomenology of CA dynamics, which includes the qualitative Wolfram classes, the existence of complexity classes, the capacity for universal computation, undecidability, and so forth.

Second, we note that the surface-level phenomenology of CA systems is remarkably similar to the surface-level phenomenology of computational systems, which includes the existence of complexity classes, computability classes, the capacity for universal computation, undecidability, etc.

Third, we conclude that the structure of the space of computations is dominated by a second-order phase transition, in terms of which the existence of—and relationships between—the surface-level features of computational systems find a relatively straightforward explanation.

### 6.1 THE FUNDAMENTAL STRUCTURE OF CA RULE SPACE

Piecing together the results of the previous sections results in a clear picture of the fundamental structure of cellular automata rule space. This structure is illustrated schematically in Figure 21.

There are two primary regimes of rules—*periodic* and *chaotic*— separated by a *transition regime*. This transition regime is *not* simply a smooth surface between the other two domains, but itself has a complicated structure. From the perspective provided by the $\lambda$ parameter, it seemed that much of the transition regime is

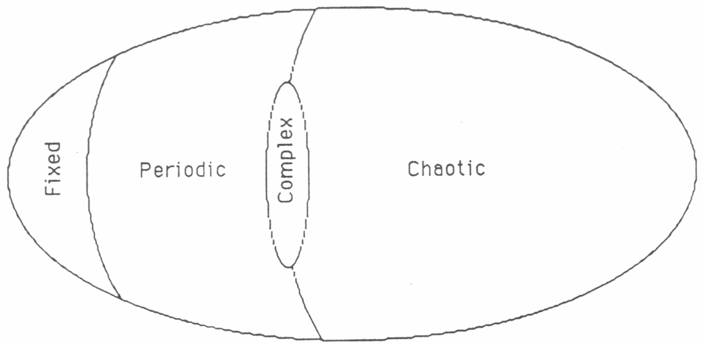


FIGURE 21 Schematic drawing of CA rule space indicating relative location of periodic, chaotic, and complex regimes.

simply a boundary between periodic and chaotic rules, containing no rules within it. Crossing the transition regime at such a boundary appeared to give rise to a discrete jump from "simple" periodic dynamics to "simple" chaotic dynamics, accompanied by a discrete jump in statistical measures of the kind usually associated with first-order transitions.

However, from a more detailed investigation[8] it is apparent that the phase transition is primarily a second-order, or *critical*, transition. Crossing this critical transition region gives rise to "complex" dynamics, accompanied by relatively "smooth" changes in statistical measures.

## 6.2 PHASE TRANSITIONS AND CA PHENOMENOLOGY

The phase-transition structure underlying CA rule spaces provides a coherent framework for organizing and explaining much of the documented phenomenology of CA dynamics.

First, this picture of a phase transition separating a domain of ordered dynamics from a domain of disordered dynamics (each of which might be further subdivided), provides a simple explanation for the existence of—and the relationships between—the four qualitative classes of CA behavior identified by Wolfram. Class IV rules are found in the transition regime separating the periodic rules (Classes I & II) from the chaotic rules (Class III). Figure 21 illustrates the way in which the Wolfram classes fit into the phase-transition picture of CA rule space.

[8]See Langton.[28]



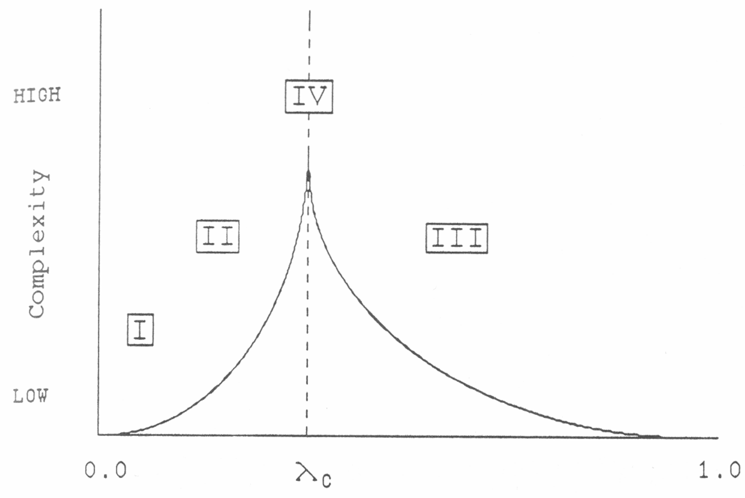


FIGURE 22 Schematic drawing of complexity versus $\lambda$ over CA rule space, showing the relationship between the Wolfram classes and the underlying phase-transition structure.

However, it is clear that the Wolfram classes, or any other scheme which attempts to partition CA rule space into a small number of "classes" or "categories," can only constitute coarse-grained approximations to the real situation. The phase portrait of CA rule space is a much more "continuous" structure with few, if any, well-defined boundaries besides the transition itself.

As one varies $\lambda$ over its full range, one sees the full spectrum of CA dynamical behaviors. Some are simple, some are complex. Varying $\lambda$ throughout its range produces a reasonably smooth progression through what appear to be two, mirror-image complexity hierarchies, one on either side of the transition regime. Approaching the transition from "below," one progresses from simple fixed-point to simple periodic behavior, and then on to longer and longer period behavior, accompanied by longer and longer transients showing more and more sensitivity to array size, until the transition regime is reached. This regime exhibits very complex behavior accompanied by extremely long transients, which show roughly exponential dependence on array size. Furthermore, finite initial configurations started in small portions of very large arrays *may or may not* exhibit transients leading to periodic behavior. The slow growth, and occasional local collapse of complex dynamics to periodic behavior, makes the ultimate outcome of a particular rule operating on a particular initial configuration impossible to predict in the general case.

As one proceeds past the transition regime into the chaotic regime, one traverses another complexity hierarchy, but in reverse order. Even just slightly past the transition regime, ultimate collapse to periodic behavior becomes extremely rare in finite arrays, and would never occur in infinite arrays. Also, transient times to typical chaotic behavior, and their dependence on array size, are shrinking rapidly with movement away from the transition regime. As one approaches the upper limit of the range for $\lambda$, CA become completely randomized after minimal transients with *no* dependence on array size. Although maximally chaotic, these dynamics are easily predictable in a statistical sense, and are therefore much simpler than the dynamics observed in the transition regime.

Thus, by varying $\lambda$, one proceeds *up* through one complexity hierarchy, from simple to complex dynamics, and then proceeds *down* through another such hierarchy, from complex to simple dynamics. At either end, behavior is simple and predictable, while in the middle, behavior can become arbitrarily complex and highly unpredictable. The entire spectrum of CA dynamical behaviors can be ordered with respect to "distance" from a critical transition point in CA rule space.

## 6.3 THE PHENOMENOLOGIES OF CA AND COMPUTATIONS

From the discussions above, it is apparent that an underlying phase-transition structure is responsible for the existence of most of the surface-level features of the diverse phenomenology of CA behaviors. Furthermore, the relationship between many of these features is explained coherently by the existence of this underlying phase-transition structure.

When we look to the diverse phenomenology of computations, we see a remarkably similar set of surface-level features as those documented above for CA, including complexity hierarchies, computability classes, arbitrarily long transients leading to unpredictability, terminating and non-terminating dynamics, universal computational capacity, and so forth.

In this section, we briefly review the phenomenology of computations and show how similar it is to the phenomenology of CA described above. In the next section, we see that the phenomenology of computation can be equally well explained by assuming a phase transition underlying the space of computations.

**6.3.1 COMPUTABILITY CLASSES** As we have seen, there are three possibilities for the ultimate outcome of the evolutions of CA. Some CA rapidly "freeze-up" into short-period behavior from any possible initial configuration. On the other hand, most CA will never become periodic, rapidly settling down to chaotic behavior instead. We can predict the ultimate dynamics of many CA with a high degree of certainty.

However, there exist some CA for which both of these ultimate dynamical outcomes are possible, and because they are typically associated with extended transients, it is effectively undecidable whether a particular CA rule operating on a particular initial configuration will ultimately lead to a periodic state or not.

Similarly, some computations halt, some do not. For most computations, we can decide whether or not they will halt. However, Turing demonstrated that this "Halting problem" is, in general, undecidable. There exist computations for which



it is *not* possible to decide whether or not they will halt when started on arbitrary inputs.

Thus, with respect to our ability to decide the ultimate behavior of both CA and computations, there are essentially three possibilities: we can determine that they will halt, we can determine that they will not halt, or we *cannot* determine whether or not they will halt.

**6.3.2 COMPLEXITY CLASSES** As we saw for CA, some CA relax to their ultimate dynamics after transients whose lengths are independent of the size of the array, while for other CA this transient time can depend exponentially on array size. This is true for CA both in the periodic and in the chaotic regime.

Similarly, for computations that *do* halt, some halt in an amount of time which is only a linear—or even a constant—function of the "size" of the input, while other computations exhibit polynomial, or even exponential, dependence on input size.[17] These different classes of functions describing the relationship between input size and "transient time" to halting constitute a complexity hierarchy, with the simplest computations at one end, and the most powerful computations at the other end.

Furthermore, by resorting to the device of "oracles,"[21] it has been determined that there is a similar complexity hierarchy for *non-halting* computations.

**6.3.3 UNIVERSALITY** We have seen that the most complex CA have the interesting property that they are arbitrarily "programmable." This means that by merely manipulating details of the initial configuration, without changing the rules of the CA itself, any computable function can be implemented

Similarly, it is known that there exist Turing Machines—representatives of the maximal class of computing devices—that can imitate any other Turing Machine merely through manipulations of the initial state of the input tape.

Thus, both CA and the more familiar realizations of computing devices can be constructed so that their global dynamics is arbitrarily sensitive to subtle details of their initial state.

## 6.4 THE "DEEP STRUCTURE" OF COMPUTATION

As is clear from the section on CA and phase transitions above, these familiar features of the phenomenology of computation have a clear and simple explanation if we assume that a critical phase transition dominates the structure of the space of computations. This structure is illustrated in Figure 23, which should be compared to Figure 22.

First, the existence of computability classes is explained by the fact that the phase transition separates the space of computations into an ordered and a disordered regime, which we refer to as the halting and the non-halting computations, respectively.

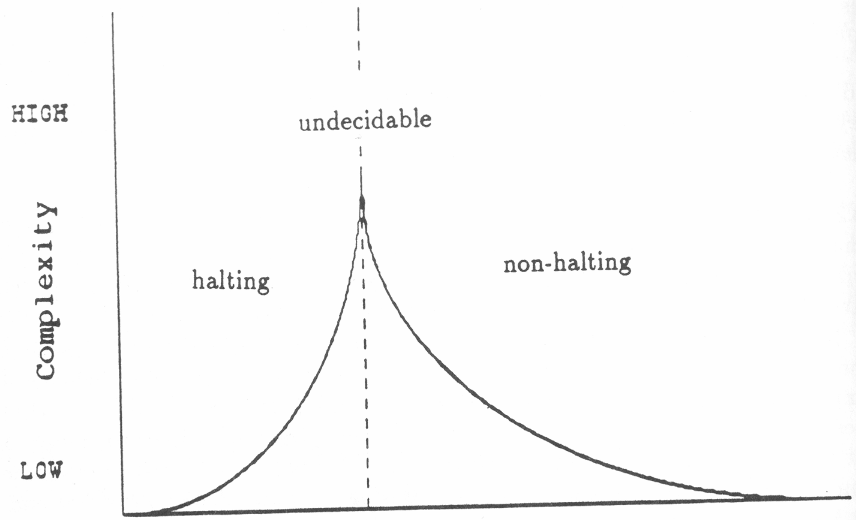


FIGURE 23 Schematic drawing of computation space, indicating relative location of halting, non-halting, and undecidable regimes.

Second, the dynamics in the vicinity of a second-order transition gives rise to the phenomenon of *critical slowing down*, and therefore the transients along the way to ultimately periodic (halting) behavior can diverge to infinity, which explains the existence of undecidability, as characterized by Turing's famous Halting Problem.

Third, the existence of complexity hierarchies is explained by the functional relationship between expected transient time, system size, and distance (with respect to some parameter) from the actual transition point observed for systems exhibiting second-order phase transitions. This functional relationship in the divergence of transient times occurs whether approaching the transition from above or below, and hence, we find that the most complex computations are to be found at a phase transition, between the halting and the non-halting domains, sandwiched in between two complexity hierarchies that diverge as one approaches the phase transition from either the halting regime or from the non-halting regime. Away from the transition regime, transient times are independent of system size, while close to the transition itself, transients can grow exponentially with system size.

Fourth, the existence of universal computation is explained by the fact that the dynamics of physical systems in the vicinity of critical transitions exhibit a divergence in their "susceptibility," that is, in their sensitivity to minute details of their internal structure and to external perturbations. This self-sensitivity in the vicinity of a critical transition is manifested in universal computers as their "programmability."



Thus, if we assume that a critical phase transition dominates the structure of the space of computations, *we expect to see exactly what we do see* when we review the phenomenology of computation.

We *expect* to see a divergence in transient lengths as one approaches the computations in the vicinity of the transition, and we expect to see this increase whether we approach these computations from either the ordered (halting) or the disordered (non-halting) regimes. We *expect* that the functional relationship between transient length, system size, and "distance" from the transition should be governed by different classes of bounding functions, increasing in complexity from constant to exponential as one approaches the computations in the vicinity of the phase transition. We *expect* to see a divergence in susceptibility (programmability) for the computations in the vicinity of the transition. We *expect* to see critical slowing down (the Halting Problem) for the computations in the vicinity of the phase transition.

Thus, the assumption that a critical transition underlies the space of computations provides a simple and straightforward explanation for the existence of, and the relationship between, many significant features of the phenomenology of computation, features which are currently known only through a loose collection of theorems, hypotheses, lemmas, corollaries, and observations.

## 7. IMPLICATIONS AND QUESTIONS

Now, returning to the main theme of this paper, what does all of this have to do with life? Simply stated, it means that we now have evidence for a *natural domain of information* in the physical world. We can now provide a tentative answer to the question posed at the beginning of this paper:

> *We expect that information processing can emerge spontaneously and come to dominate the dynamics of a physical system in the vicinity of a critical phase transition.*

As the origin of life is intimately associated with the spontaneous emergence of a dynamics of information in the physical world, we suggest that the origin of life occurred when some physico-chemical process underwent a critical phase transition in the early history of the Earth.

In the following sections, we will review what I believe to be some of the implications of the existence of a fundamental connection between computation and phase transitions. What follows is, of necessity, more speculative than the previous sections, as it anticipates future work, rather than describing work already accomplished. However, a reasonable case can be made for everything that follows.

## 7.1 PHYSICS AND COMPUTATION

The results of this research point to a fundamental equivalence between the dynamics of phase transitions and the dynamics of information processing. Such a connection is a two-edged sword. With one edge, we can apply what we know about the physics of phase transitions to further our understanding of computation. With the other edge, we can apply what we know about computation to further our understanding of phase transitions.

Thus, the establishment of such a connection will bring us fundamentally new insights into the nature of computation. Some of the phase-transition phenomena mentioned above as analogs of computational phenomena are well understood in the context of such bodies of theory as thermodynamics, statistical mechanics, and renormalization group theory. By understanding computation as a special case of phase-transition phenomena, much of the apparatus of the above bodies of theory can be brought to bear on problems in the theory of computation.

Likewise, the apparatus of the theory of computation can be brought to bear on problems in the theory of phase transitions. In particular, it may become clear that certain aspects of phase transitions are not treatable by *any* theory in principle, because they effectively involve undecidable problems. This connection already tells us something about the phenomenon of critical slowing down: it is the physical manifestation of the fact that the system is engaged in "solving" an intractable problem. A material near its critical transition point between the liquid and gas states must, in effect, come to a global decision about whether it will settle down to a liquid or to a gas. This sounds almost anthropomorphic, but the results reported here suggest that we must think of such systems as effectively *computing* their way to a minimum energy state.

It also lends support to the Church-Turing hypothesis that *no* system will be found that computes a wider class of functions than a Turing machine. The fact that physical systems having to "decide" between two qualitatively different physical states can take arbitrarily long to make the "decision"—the phenomenon of critical slowing down—suggests that physical systems in general are bound by the same *in principle* limitations as computing devices.

In fact, it is perhaps misleading to claim that the phenomenology of phase transitions explains the phenomenology of computation when, in fact, one could equally well claim that the phenomenology of computation explains the phenomenology of phase transitions. On the basis of the evidence presented here, we could equally well say that the phenomenon of critical slowing down is simply a physical manifestation of the halting problem.

What all this is really telling us is that these two phenomenologies are not "two" at all but rather "one." We are observing one and the same phenomenon reflected in two very different classes of systems and their associated bodies of theory.



## 7.2 SOLIDS, FLUIDS, AND DYNAMICS

With CA, and other spatially extended dynamical systems, we are experimenting with "artificial matter" (or "programmable matter" in Toffoli's words[45]). That is, we are experimenting with "material" for which we have precise control over the behaviors of the individual "atoms" or "molecules" of which it is constituted.

It is somewhat surprising that, despite the many different varieties of atoms and molecules that constitute "real" materials, almost every known substance comes in one of three flavors: solid, liquid, or gas. As it is possible to continuously transform liquids into gases and *vice versa* without passing through a phase transition, they are taken to constitute a single, more general phase of matter: *fluids*. Therefore, there are really just *two* fundamental phases of matter—solids and fluids—and so we should not be surprised to find two similar fundamental phases in our "artificial" materials, despite the large number of ways that we can put them together. The important point here is that solids and fluids are *dynamical*, rather than merely *material*, categories.

We know solids and fluids primarily as states of matter, rather than as universal classes of dynamical behavior, because up until quite recently, everything that exhibited dynamical behavior was fundamentally material in constitution. Now, however, with the availability of computers, we are able to experiment with dynamical behaviors *per se*, abstracted from any particular material substrate. What we find is that, despite having abandoned the material basis of solids and fluids, *we are nonetheless left in possession of solid and fluid dynamics!* Thus, we are safe in assuming that these fundamental classes of dynamical behavior do not inhere in material *per se*, but rather in the way in which the material is organized.

The most important point, however, is that these two universality classes of dynamical behavior are separated by a *phase transition*. As we have seen, the dynamics of systems operating near this phase transition provides the basis of support for embedded computation. Thus, a third category of dynamical behaviors exists in which materials—or more broadly, dynamical systems in general—can make use of an intrinsic computational capacity to avoid either of the two primary categories of dynamical behavior by maintaining themselves on indefinitely extended transients.

Since computers and computations are specific instances of material and dynamical systems respectively, they are also ultimately bound by these same universality classes. Therefore, computer "hardware" can behave like a solid, like a fluid, or like something in between.

An interesting open question to pursue here is whether one can resolve the fluid phase of CA (and other dynamical systems) even further into liquid and vapor phases. The careful reader may have noticed that I have consistently used the phrase "in the vicinity of a second-order phase transition" rather than the phrase "*at* a second-order phase transition." This is in part due to the fact that it is hard to determine whether or not one is precisely "at" a critical point when working with finite systems. But it is also due to the fact that in most of the experiments on CA leading to these results, the regime exhibiting the most complex dynamics appears to be just slightly below (to the ordered side) of the critical transition itself. Most physical systems exhibit a liquid phase or the "ordered side" of their critical point, and liquids exhibit extremely complex dynamics.

It is somewhat surprising that liquids are *very* poorly understood. Although the structures of both solids and gases are well studied, it has proven difficult in practice to characterize the structure of liquids. As might be expected from their location between the solid and the gas phases, liquids exhibit much more complicated behaviors than either of the other two phases. Recent work by Stanley[43] suggests that the hydrogen-bond network of liquid water exhibits very complex dynamics, and super-cooled water can apparently exhibit critical dynamics.

It may be appropriate to view liquids as constituting a broadened phase transition in the phase portrait of a material. Liquids occupy just the right spot "at the edge of chaos" in the phase portrait of many materials to be candidates for further investigation along the lines of the results reported in this paper. The complex transition regime might be co-extensive with the liquid regime for many real and "artificial" materials.

## 7.3 LIFE, INTELLIGENCE, AND EVOLUTION

This generalization of the solid, liquid, and vapor phases of matter to universal classes of behavior for dynamical systems in general, has important consequences for our understanding of the kinds of behaviors achievable by computer "hardware." It has generally been thought that computer hardware and biological "wetware" are fundamentally different kinds of "stuff," and that, because of this fundamental difference, computer hardware could never achieve some of the more exotic dynamical behaviors exhibited by wetware, such as life and intelligence. However, if it is properly understood that hardness, wetness, or gaseousness are properties of the organization of matter, rather than properties of the matter itself, then it is only a matter of organization to turn "hardware" into "wetware" and, ultimately, for "hardware" to achieve everything that has been achieved by wetware, and more.

In addition, this perspective implies that a material system operating in the vicinity of a phase transition can behave like a computer. This has important consequences for our understanding of the kinds of behaviors achievable by "wetware."

Life is clearly founded on a capacity to sense, process, and act on information. The primary macromolecules of life are "informational" molecules, and these molecules are engaged in a finely tuned dynamics of information processing which we have only just begun to understand.

It has long been an open question how such a dynamics of information could have gotten started within the pre-biotic soup, and how it could have refined itself and evolved through time to produce such magnificent information processors as human brains. The results of this research suggest a new set of answers to these questions.

We now know that a dynamics of information can spontaneously emerge in physical systems near a critical phase transition. The results of this paper suggest the possibility that the information dynamics which gave rise to life came into



existence when global or local conditions brought some medium—perhaps $H_2O$, perhaps some other material—through a critical phase transition. As we have seen in CA, the dynamics in the vicinity of a critical phase transition support a complex dynamics of metastable structures, which is very much of a kind with the complex dynamics of metastable structures that characterizes life. Furthermore, we also know that, in CA at least, this dynamical regime can support arbitrarily complicated information processing. Thus, this regime is a good candidate for the origin of the kind of information dynamics that we would consider to be synonymous with the origin of life.

Now, Nature could not have been so beneficent as to have maintained this medium near a critical phase transition for very long. This means that the nascent information dynamics must have gained control over some parameters that allowed them to maintain *local* conditions near the phase transition while *global* conditions drifted away from the phase transition. Of course, there must be many parameters that could push such systems away from their vital transition point, and many of these probably varied widely, destroying a large proportion of these early information-processing systems. Evolution can be viewed as the process of gaining control over more and more "parameters" affecting a system's relationship to the vital phase transition.

Living systems can perhaps be characterized as systems that dynamically *avoid* attractors. The periodic regime is characterized by limit-cycle or fixed-point attractors, while the chaotic regime is characterized by strange attractors, typically of very high dimension. Living systems need to avoid either of these ultimate outcomes, and must have learned to steer a delicate course between too much order and too much chaos—the *Scylla* and *Charybdis* of dynamical systems.

They apparently have done so by learning to maintain themselves on extended transients—i.e., *by learning to maintain themselves near a "critical" transition.* Once such systems emerged near a critical transition, evolution seems to have discovered the natural information-processing capacity inherent in near-critical dynamics, and to have taken advantage of it to further the ability of such systems to maintain themselves on essentially open-ended transients.

Of course, climbing out of one attractor just pushes the problem back to a higher-dimensional phase space, in which the system is again in the basin of some attractor. It is therefore possible to view evolution as a repeated iteration of the process whereby a system climbs out of one attractor into a higher-dimensional phase space, only to find itself in the basin of a higher-dimensional attractor, a process that gives evolutionary significance to the phrase "out of the frying pan, into the fire!"

In the context of the work of Stanley, described in the previous section, it is interesting to consider the possibility that the dynamics of information which eventually lead to the origin of life may have emerged as the structural dynamics of liquid water. Theories of the origin of life have assumed that life originated in the dynamics of molecules *embedded* in liquid water. However, rather than liquid water having merely provided the "nursery" for the origin of life among molecules embedded within it, life may have originated in the dynamics of water itself. The dynamics may then have eventually spread to the embedded molecules, setting the stage for a "genetic takeover" of the kind proposed by Cairns-Smith,[8] in which the induced dynamics of the embedded molecules took over from the original dynamics of liquid water. Or, perhaps the natural dynamics of liquid water still plays a "vital" role in modern life.

There is ample evidence in living cells to support an intimate connection between phase transitions and life. Many of the processes and structures found in living cells are being maintained at or near phase transitions. Examples include the lipid membrane, which is kept in the vicinity of a sol-gel transition; the cytoskeleton, in which the ends of the microtubules are held at the point between growth and dissolution; and the naturation and de-naturation (zipping and unzipping) of the complementary strands of DNA.

In the case of intelligence, there is also qualitative evidence for this phase transition view in the dynamics of the brain. It is vital that the brain be kept very near to 98.6°F in order to work properly. We've all experienced the chaotic nature of our thinking processes when we have a fever. Some have experienced the seizures (periodic dynamics) that accompany hypothermia, when the brain gets too cold. On the temperature scale, clearly, the brain operates in a very narrow regime between periodic and chaotic dynamics, and a great amount of physiological machinery has evolved to keep it at this critical point. Our mental capabilities are apparently only possible in the vicinity of this phase transition between periodic and chaotic neural dynamics.

There is also evidence that evolutionary dynamics brings populations to "the edge of chaos." Eldredge and Gould claim that the fossil record exhibits *punctuated equilibria*.[14] That is, the fossil record seems to indicate that long periods of evolutionary stasis are irregularly interrupted by periods of chaotic and rapid evolutionary change. Evidence from recent computer experiments indicates that this phenomenon may be generic for evolutionary processes.[33] Many nonlinear dynamical systems exhibit a remarkably similar phenomenon known as "intermittency," in which long periods of regular periodic behavior are interrupted irregularly by bursts of rapidly fluctuating chaotic behavior. Most importantly, intermittency is generically observed in these systems when they are in the vicinity of a transition from periodic to chaotic behavior. Thus, punctuated equilibria in the fossil record can be viewed as evidence for the proposition that biological evolution has maintained populations in a state of near-critical dynamics.

## 7.4 RELATED WORK

There are other lines of research that have intimate connections with this work.

As mentioned earlier, Jim Crutchfield has observed a similar increase in effective computational complexity in the vicinity of phase transitions in continuous dynamical systems.[11]

Vichniac, Tamayo, and Hartman[46] discovered that the Wolfram classes could be recovered by varying the frequency of two simple rules in an inhomogeneous cellular automaton. They also suggested a relation between critical slowing down and the halting problem.



Norman Packard and Wentian Li have mapped out the space of "elementary" $K = 2$, $N = 3$, one-dimensional CA fairly completely, using a parameterization scheme similar to $\lambda$.[31]

Packard has also performed an interesting series of experiments in which he "adapts" CA rules by selecting for certain properties of the global dynamics.[39] He finds that an initially random population of rules will migrate towards the phase-transition region. His interpretation of this phenomenon is that it is easier to find rules that will *compute* the desired behavior—by making use of a general computational capacity—than it is to find rules that are "hard-wired" to produce *only* the desired behavior.

Stuart Kauffman has investigated a class of dynamical systems known as *boolean nets*,[22,23] in which he finds a phase transition between ordered and disordered dynamics as a function of the *connectivity* of the network. These nets also exhibit phase transitions as a function of an "internal homogeneity parameter," a parameter very much like $\lambda$, which controls the number of 1's in the boolean functions found at each node in the network.[47] Kauffman is currently working on a set of experiments to see if evolving populations of these boolean nets will converge on the transition regime in their parameter space (see the contribution by Kauffman and Johnsen in these proceedings.)

Harold McIntosh[36] has applied the mean-field approach of Gutowitz[19,20] and suggests that the Wolfram classes can be distinguished on the basis of simple features of the mean field theory curves. These simple features clearly locate Class IV (complex rules) at the transition between Class II (periodic rules) and Class III (chaotic rules.)

Bill Wootters has applied mean-field theory to explain the $\lambda$ parameter results, and has been able to reproduce many of the features of Figure 6.[52]

The computer-optimization procedure known as *simulated annealing*[24] has a strong connection with this work. Annealing schedules call for extended stays in the vicinity of the freezing point. This is partially due to the phenomenon of critical slowing down: it takes longer for the systems to relax near phase transitions. However, we have seen that slowing down is intimately associated with complex computational dynamics, and it is interesting that the transition is the very point at which we expect information processing to emerge spontaneously within the system being annealed. This suggests that the real reason for hovering in the vicinity of the freezing point—and the reason that it takes longer to relax there—is that even the simple act of relaxation in this regime is computationally complex: it cannot simply relax, it must *compute* its way to a minimum energy state. The system at the freezing point is effectively caught up in running an embedded computation.

The concepts reported here clearly have some relation to the notion of self-organized criticality.[3] Bak et al. have proposed that a number of systems, including some CA, exhibit a magical capacity to organize themselves towards a critical state, *in the absence of any parameter tuning*. Bak has suggested that Conway's game of LIFE is a self-organized critical system, although he does not bring LIFE's computational capacity into the discussion.[4] Bennett[5] has recently demonstrated that LIFE is generically *sub-critical*, but suggests that under a small set of initial conditions, LIFE might be *super-critical*.

Finally, Andy Wuensche and Mike Lesser have recently produced an atlas of basin of attraction fields for simple one-dimensional CA.[53] These basin fields are constituted of directed graphs on the state-spaces of CA. When organized according to the $\lambda$ parameter, or according to a related parameter $Z$ which they define, one can observe that the phase transition in CA dynamics is associated with a phase transition in the associated state-space graphs—a transition analogous to the emergence of a "giant component" in random graphs. Fixed-point and periodic CA are associated with a high degree of convergence in the state space, while chaotic CA are associated with very low convergence. Complex CA have irregular and extremely long transient trees, exhibiting a low but positive degree of convergence. Their parameter $Z$ is a global measure of the overall degree of convergence in state-space, and is probably a more accurate predictor of the phase transition than $\lambda$.

Thus, there is a growing body of experimental and theoretical evidence pointing to a fundamental association between complex dynamics, computational capacity, and phase transitions in both abstract-formal and concrete-physical systems, organic as well as inorganic.

## 8. CONCLUSION

In this paper, we have attempted to uncover the conditions under which a complex dynamics of information processing can emerge spontaneously and come to dominate the dynamics of a physical system. This has lead us to the observation that a second-order, or *critical*, phase transition underlies the space of a number of classes of dynamical systems, including cellular automata (CA) and computations, and that sytems in the vicinity of such a phase transition can support arbitrary information processing.

We have concluded that phase transitions figure "critically" in the origin and evolution of life and intelligence.

## ACKNOWLEDGMENTS

I would like to thank the following friends and colleagues for stimulating discussions on the topics discussed here: Richard Bagley, Jim Crutchfield, Doyne Farmer, Stephanie Forrest, Walter Fontana, Howard Gutowitz, Hyman Hartman, Stuart Kauffman, Wentian Li, Norman Packard, Steen Rasmussen, Rob Shaw, Creighton Walker, and Bill Wootters.



## REFERENCES

1. Atkins, P. W. *Physical Chemistry*. San Francisco: W. H. Freeman, 1978.
2. Atkins, P. W. *The Second Law*. New York: W. H. Freeman, 1984.
3. Bak, P., C. Tang, C., and K. Wiesenfeld. "Self-Organized Criticality." *Phys. Rev. A* **38(1)** (1988): 364–374.
4. Bak, P., K. Chen, and M. Creutz. *Nature* **342** (1989): 780–782.
5. Bennett, C., and M. S. Bourzutschky. "Non-Criticality of the Cellular Automaton 'Life.'" *Nature* (1990).
6. Berlekamp, E., and J. H. Conway, and R. Guy. *Winning Ways for Your Mathematical Plays*. New York: Academic Press, 1982.
7. Burks, A. W. *Essays on Cellular Automata*. Urbana: University of Illinois Press, 1970.
8. Cairns-Smith, A. G. *Genetic Takeover and the Mineral Origins of Life*. Cambridge: Cambridge University Press, 1982.
9. Chaitin, G. *J. Assoc. Comput. Mach.* **13** (1966): 145.
10. Codd, E. F. *Cellular Automata*. New York: Academic Press, 1968.
11. Crutchfield, J. P., and K. Young. "Computation at the Onset of Chaos." *Complexity, Entropy, and Physics of Information*, edited by W. H. Zurek, 223–270. Santa Fe Institute Studies in the Sciences of Complexity, Proc. Vol. VIII. Redwood City, CA: Addison-Wesley, 1990.
12. Culik, K., and S. Yu. "Undecidability of CA Classification Schemes." *Complex Systems* **2** (1988): 177–190.
13. Derrida, B., and Y. Pomeau. "Random Networks of Automata: A Simple Annealed Approximation." *Europhys. Lett.* **1(2)** (1986): 45–49.
14. Eldredge, N., and S. J. Gould. "Punctuated Equilibria: An Alternative to Phyletic Gradualism." *Models in Paleobiology*, edited by T. J. M. Schopf, 82–115. Freeman Cooper, 1972.
15. Fredkin, E., and T. Toffoli. "Conservative Logic." *Int'l J. Theor. Phys.* **21** (1982): 219–253.
16. Gardner, M. "Mathematical Games: The Fantastic Combinations of John Conway's New Solitaire Game 'Life.'" *Sci. Am.* **223(4)** (1970): 120–123.
17. Garey, M. R., and D. S. Johnson. *Computers and Intractability*. San Francisco: W. H. Freeman, 1979.
18. Gatlin, L. L. *Information Theory and the Living System*. New York: Columbia University Press, 1972.
19. Gutowitz, H. A., B. K. Knight, and J. D. Victor. "Local Structure Theory for Cellular Automata." *Physica D* **48** (1987): 18.
20. Gutowitz, H. A. "A Hierarchical Classification of Cellular Automata." *Proceedings of the 1989 Cellular Automata Workshop*, edited by H. A. Gutowitz. Amsterdam: North Holland, 1990.
21. Hopcroft, J. E., and J. D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Menlo Park, CA: Addison-Wesley, 1979.

22. Kauffman, S. A. "Metabolic Stability and Epigenesis in Randomly Constructed Genetic Nets." *J. Theor. Biol.* **22** (1969): 437–467.
23. Kauffman, S. A. "Emergent Properties in Random Complex Automata." *Physica D* **10** (1984).
24. Kirkpatrick, S., C. D. Gelatt, and M. P. Vecchi. "Optimization by Simulated Annealing." *Science* **220(4598)** (1983): 671–680.
25. Kolmogorov, A. N. *Prob. Inf. Transm.* **1(1)** (1965).
26. Langton, C. G. "Studying Artificial Life with Cellular Automata." *Physica D* **22** (1986): 120–149.
27. Langton, C. G. "Computation at the Edge of Chaos." *Physica D* **42** (1990): 12–37.
28. Langton, C. G. "Computation at the Edge of Chaos: Phase-Transitions and Emergent Computation." Ph.D. thesis, University of Michigan, 1991.
29. Langton, C. G., C. Taylor, J. D. Farmer, and S. Rasmussen, eds. This volume.
30. Li, W. "Analyzing Complex Systems." Columbia University, 1989.
31. Li, W., and N. Packard. "Structure of Elementary Cellular Automata Rule-Space." *Complex Systems* (1990).
32. Lindgren, K., and M. G. Nordahl. "Universal Computation in Simple One-Dimensional Cellular Automata." *Complex Systems* **4** (1990): 299–318.
33. Lindgren, K. "Evolutionary Phenomena in Simple Dynamics." This volume.
34. Mandelbrot, B. B. *Fractals: Form, Chance, and Dimension.* San Francisco: W.H. Freeman, 1977.
35. Mayr, E. *The Growth of Biological Thought.* Cambridge: Harvard University Press, 1982.
36. McIntosh, H. V. "Class IV Cellular Automata and a Good LIFE." *Proceedings of the 1989 Cellular Automata Workshop*, edited by H. A. Gutowitz. Amsterdam: North Holland, 1990.
37. von Neumann, J. *Theory of Self-Reproducing Automata*, edited and completed by A. W. Burks. Urbana: University of Illinois Press, 1966.
38. Packard, N. H., and S. Wolfram. "Two-Dimensional Cellular Automata." *J. Stat. Phys.* **38** (1985): 901.
39. Packard, N. H. "Adaptation Toward the Edge of Chaos." *Complexity in Biological Modelling*, edited by S. Kelso and M. Shlesinger. 1988.
40. Reichl, L. E. *A Modern Course in Statistical Physics.* Austin: University of Texas Press, 1980.
41. Smith III, A. R. "Simple Computation-Universal Cellular Spaces." *JACM* **18(3)** (1971): 339–353.
42. Stanley, H. E. *Introduction to Phase Transitions and Critical Phenomena.* Oxford: Oxford University Press, 1971.
43. Sciortino, P. H. P., H. E. Stanley, and S. Havlin. "Lifetime of the Bond Network and Gel-Like Anomalies in Supercooled Water." *Phys. Rev. Lett.* **64(14)** (1990): 1686–1689.
44. Stauffer, D. *Introduction to Percolation Theory.* Philadelphia, PA: Taylor and Francis, 1985.





45. Toffoli, T., and N. Margolus. *Cellular Automata Machines.* Cambridge MA: MIT Press, 1987.
46. Vichniac, G. Y., P. Tamayo, and H. Hartman. "Annealed and Quenched Inhomogeneous Cellular Automata." *J. Stat. Phys.* **45** (1986): 875–883.
47. Walker, C. Personal communication.
48. Wolfram, S. "Stastical Mechanics of Cellular Automata." *Rev. Mod. Phys.* **55** (1983): 601–644.
49. Wolfram, S. "Universality and Complexity in Cellular Automata." *Physica D* **10** (1984): 1–35.
50. Wolfram, S. "Twenty Problems in the Theory of Cellular Automata." *Physica Scripta* (1985).
51. Wolfram, S. *Theory and Applications of Cellular Automata.* Singapore: World Scientific, 1986.
52. Wootters, W. T., and C. G. Langton. "Is There a Sharp Phase Transition for Deterministic Cellular Automata?" *Physica D* **45** (1990): 95–104.
53. Wuensche, A., and M. Lesser. *An Atlas of Basin of Attraction Fields of One-Dimensional Cellular Automata* (working title). Santa Fe Institute Studies in the Sciences of Complexity, Reference Volume I. Redwood City, CA: Addison-Wesley, 1992.