

LEGIBILITY NOTICE

A major purpose of the Technical Information Center is to provide the broadest dissemination possible of information contained in DOE's Research and Development Reports to business, industry, the academic community, and federal, state and local governments.

Although a small portion of this report is not reproducible, it is being made available to expedite the availability of information on the research discussed herein.

APR 05 1990

Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36

LA-UR--90-878

DE90 008914

TITLE: Is There a Sharp Phase Transition
for Deterministic Cellular Automata?

AUTHOR(S): William K. Wootters, Santa Fe Institute,
CNLS/T-13 Los Alamos National Laboratory
Dept. of Physics, Williams College
Chris G. Langton, T-13 Los Alamos National Laboratory

SUBMITTED TO: Physica D

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution or to allow others to do so, for U.S. Government purposes.

The Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy.

 **Los Alamos** Los Alamos National Laboratory
Los Alamos, New Mexico 87545

Is There a Sharp Phase Transition for Deterministic Cellular Automata?

William K. Wootters^(1,2,3) and Chris Langton⁽²⁾

⁽¹⁾Santa Fe Institute, 1120 Canyon Road, Santa Fe, NM 87501

⁽²⁾T-13 and Center for Nonlinear Studies
Los Alamos National Laboratory, Los Alamos, NM 87545

⁽³⁾Department of Physics, Williams College, Williamstown, MA 01267

Abstract

Previous work has suggested that there is a kind of phase transition between deterministic automata exhibiting periodic behavior and those exhibiting chaotic behavior. However, unlike the usual phase transitions of physics, this transition takes place over a range of values of the parameter rather than at a specific value. The present paper asks whether the transition can be made sharp, either by taking the limit of an infinitely large rule table, or by changing the parameter in terms of which the space of automata is explored. We find strong evidence that, for the class of automata we consider, the transition does become sharp in the limit of an infinite number of symbols, the size of the neighborhood being held fixed. Our work also suggests an alternative parameter in terms of which it is likely that the transition will become fairly sharp even if one does not increase the number of symbols. In the course of our analysis, we find that mean field theory, which is our main tool, gives surprisingly good predictions of the statistical properties of the class of automata we consider.

1 Introduction

Of the four Wolfram classes [1] of deterministic cellular automata--homogeneous, periodic, chaotic, and complex--those exhibiting complex behavior are for many purposes the most interesting, but they also seem to be the most rare, particularly when the rule table for the automaton is large [2]. Thus, when the rule table is large one can think crudely of the set of all rules as being divided into two large classes: those leading to periodic behavior (we count a static homogeneous automaton as trivially periodic), and those leading to chaotic behavior.

Langton [3,4,5] and Li and Packard [6,7] have shown that if one uses an appropriate parameter, which we call λ , to characterize cellular automaton rules, then one can see a kind of phase transition between these two modes of behavior as the value of the parameter is varied: For small values of λ the behavior is typically periodic; for large values it is typically chaotic. However, there is not a unique value of λ at which the transition occurs, but rather a range of possible transition values. In this respect the transition differs from the usual phase transitions of physics, which occur at definite values of, say, temperature or magnetic field strength. The present paper asks whether in some appropriate limit--we consider particularly the limits of large neighborhood size and large number of symbols--or with a different choice of parameter, the value at which the transition occurs becomes unique. Our results indicate that the transition does become sharp in the limit of an infinite number of symbols. We also propose a manageable experiment in which one might hope to see a reasonably sharp transition with a small set of symbols.

The question of the sharpness of the transition is interesting for a number of reasons. First, there is a good deal of evidence that automata exhibiting complex behavior are located near the transition region. If there are situations in which the transition is sharp, then in those situations one could conceivably be better able to find complex automata, because one would know precisely where to look. Also, Packard has obtained data suggesting that cellular automaton rules, evolving under selective pressure towards performing a particular computational task, tend to gravitate to the transition region [8]. It would be interesting to see, in such an evolution, if the final evolved rules continue to be confined to the transition region as the latter becomes narrower and narrower. Finally, there are some intriguing parallels between the theory of computation and the theory of physical phase transitions [4,5], and these parallels would be given added significance if one could find in deterministic cellular automata the sort of sharp transition that one finds in physics.

The approach taken in this paper is the following. The location of the transition point is a difficult quantity to treat theoretically. But there is a more accessible quantity, the single-site entropy, which also exhibits a range of values and which is related to the location of the transition point. We study the entropy in detail, asking whether for fixed λ it approaches a unique value in the limit of large neighborhood size or large number of symbols. If it does, then it seems likely that the location of the transition point becomes unique in the same limit. We supplement our analysis of finite-state automata with experimental results bearing particularly on the case of an infinite number of symbols. Throughout this paper, we focus exclusively on a specific set of data for a particular class of automata. Let us therefore begin by describing this class.

2 Our Automata

The automata we consider are two-dimensional, with a square lattice, and with a five-cell neighborhood shaped like a "+". Each site can be in one of eight possible states; in other words, each site can be occupied by one of eight possible symbols, say, the digits 0 through 7. The symbol 0 plays a special role: if a neighborhood consists entirely of zeros, then its central site is always assigned the value zero in the next step. The rules are required to be invariant under rotations. That is, if a certain neighborhood configuration sends its central site to the state x , then any rotated version of that configuration must also send its central site to x . Except for these restrictions, all rules are allowed. The lattices actually used to generate the data were 64 units on a side, with periodic boundary conditions.

The rule table can be pictured as a long column of cross-shaped neighborhoods, with each possible neighborhood configuration represented exactly once. To the right of each configuration one places the symbol to which that particular configuration is mapped. The parameter λ is defined as the fraction of neighborhood configurations that are mapped to something other than zero. Thus a rule with very small λ has a tendency to generate many zeros, and ultimately such an automaton is likely to end up in the quiescent state consisting of nothing but zeros. As λ moves away from zero, one has the possibility of more interesting behavior.

We will be looking particularly at the dependence of entropy on λ . The data was generated as follows: Fix a value of λ , and choose at random a rule table with that particular λ -value. For this rule table, start with a random initial configuration, and

without taking any data, let the automaton run for 500 time steps to give it a chance to reach its asymptotic behavior. Then let it continue to run (for 1000 steps), and count the frequencies of occurrence of the eight possible symbols. From these frequencies p_s , compute the entropy of the automaton, defined as

$$H = - \sum_{s=0}^7 p_s \log p_s . \quad (1)$$

where the logarithm is base 2. Then plot the point (λ, H) . Now do this one hundred times for each value of λ , with a different randomly chosen rule table each time. Finally, repeat the whole process for many different values of λ ranging from zero to $7/8$. ($\lambda=7/8$ is the value for which all symbols are equally represented in the rule table. For larger values one does not expect particularly interesting behavior.) The result is the graph shown in Fig. 1.

Note that in the graph, the automata can be grouped loosely into two categories. There is a band of low-entropy automata running along the bottom of the graph, and most of the other automata have entropy above a fairly well defined line at 0.84 bits. Roughly speaking, the low-entropy automata are typically periodic, and those in the high-entropy category are typically chaotic. If one starts with the $\lambda=0$ automaton and slowly adds more nonzero elements to the rule table, one usually finds the following: The entropy remains near zero as λ increases, and then at some point it abruptly jumps to a value greater than 0.84. Occasionally the entropy jumps to an intermediate value before going above 0.84, and these intermediate automata are represented by the scatter of points lying between the two larger clumps. The value of λ at which the jump occurs varies from one trial to the next.

Thus there is a certain sense in which the transition is sudden--in a given trial it is sudden--but in another sense the transition is spread over a range of λ -values. Our question is whether this range can be made to shrink to a single value.

It is not clear what will happen to the graph of Fig. 1 as the size of the neighborhood (n) or the number of symbols (k) grows. In particular, it is not clear whether the spread in entropy for each value of λ will shrink or will stay essentially constant. In a certain sense, n and k are already quite large: there are 8^5 entries in the rule table, and the number of possible rule tables is roughly equal to $8(8^{5/4})$. (The requirement of rotational invariance reduces the number of choices by a factor of about 4.) Thus it is conceivable that we are already close to the infinite-rule-table limit.

The next four sections are devoted to accounting for the observed spread in entropy and determining how it is likely to depend on n and k .

3 Simple Mean Field Theory

Before trying to account for the spread, let us illustrate our method by finding a simple theoretical estimate of the average entropy as a function of λ .

Let v_t be the density of zeros in the pattern after t time steps. We can also think of v_t as the probability that a randomly chosen site will have a zero in it. We now compute v_{t+1} as follows. For a given site, there are two ways in which that site can get a zero in the next step: Its whole neighborhood could be filled with zeros, in which case the site gets a zero with probability 1, or its neighborhood could contain at least one nonzero symbol, in which case the site will get a zero with probability $(1-\lambda)$. We assume that the probability of the neighborhood being all zeros at time t is v_t^5 . Thus the probability that the site will get a zero in the next step is

$$v_{t+1} = v_t^5 + (1-v_t^5)(1-\lambda). \quad (2)$$

Once the automaton has reached its steady state, the density v should not change, so v should satisfy the equation

$$v = v^5 + (1-v^5)(1-\lambda). \quad (3)$$

For $\lambda \leq 0.2$ the only stable solution of this equation is $v=1$ (all zeros), and for $\lambda > 0.2$ there is again exactly one stable solution, which is less than 1. The predicted density of nonzero symbols, $1-v$, is plotted as a function of λ in Fig. 2. To estimate the entropy, we simply assume that all seven nonzero symbols appear equally often, so that the entropy is related to v by the equation

$$H = - \left[v \log v + 7 \left(\frac{1-v}{7} \right) \log \left(\frac{1-v}{7} \right) \right]. \quad (4)$$

This estimate of the average entropy is plotted against λ in Fig. 3. Notice that the curve does go roughly through the middle of the actual data. But to account for the spread in the data, we clearly need to be more sophisticated.

4 More Sophisticated Mean Field Theory

In the above calculation, the rule table was characterized by the value of a single parameter, λ . Following Wolfram [9] and Gutowitz et al. [10], we now make a more refined characterization of the rule table. Let b_j be defined as follows: For a given rule table, consider all the neighborhood configurations containing exactly j zeros; b_j is the fraction of these which are mapped to zero. In effect we are breaking up the rule table into six sections ($j=0$ to 5), such that in each section all the neighborhood configurations have the same number of zeros. The quantity $1-b_j$ is like a "local" λ for the j^{th} section. Once one knows the set $\{b_j\}$, one can easily compute the value of λ for the whole rule table, but the converse is not true. A given value of λ is consistent with many possible sets $\{b_j\}$.

Our strategy now is simple. Knowing that for each value of λ the rules were chosen at random, we can find the distribution of the b_j 's for each λ . For each set $\{b_j\}$ we can compute the mean-field-predicted value of the entropy. Because there are many sets $\{b_j\}$ for each value of λ , there will be a range of entropy values for each λ . Our question is whether this range agrees with the observed range of values.

First let us write down the mean-field equation for the density v of zeros, using the more refined estimates of probabilities provided by the b_j 's. We write the equation in terms of arbitrary values of n and k , since we are ultimately interested in knowing how things depend on these quantities. For a given site, the probability that its neighborhood has exactly j zeros is equal to $v^j(1-v)^{n-j} \binom{n}{j}$. (Here we make the usual mean field approximation that the probabilities for different sites are independent.) Given that the neighborhood has exactly j zeros, the probability that the site will get a zero is b_j . Thus the equation for a stationary value of v is

$$v = v^n + \sum_{j=0}^{n-1} v^j (1-v)^{n-j} \binom{n}{j} b_j. \quad (5)$$

Here we have used the fact that b_n is equal to unity; that is, a neighborhood containing nothing but zeros is certain to yield a zero. The contribution from such a neighborhood is the first term on the right-hand side. (We use mean field theory here only to find the stationary density. For an example of what can be done with the full dynamical mean field theory, see [11].)

It is now a matter of mathematics to find the distribution of the b_j 's for each λ , and from that to find the distribution of values of v , and finally to find the distribution of values of the entropy. Let us begin by finding the distribution of the b_j 's.

In the actual experiment, the rule tables were constructed by proceeding through all rotationally inequivalent neighborhood configurations, and assigning to each one the symbol zero with probability $1-\lambda$, and a nonzero symbol with probability λ . There are approximately $8^{5/4}$ independent choices in the construction of such a rule table. Therefore, if one constructs many rule tables in this way, the standard deviation in the actual value of λ is given approximately by $\Delta\lambda = \sqrt{\lambda(1-\lambda)/N}$, where $N \approx 8^{5/4}$. This is the usual formula for the standard deviation of a binomial distribution. For $\lambda=1/2$, the value of $\Delta\lambda$ is 0.006, which is much too small to explain the observed spread in entropy. However, the spread in the b_j 's is significantly larger, since each b_j is determined by fewer choices. Because of the rotational invariance, finding the standard deviation of each b_j requires some careful counting, but for a fixed neighborhood size and fixed number of symbols it is a straightforward matter to do the necessary arithmetic. It is convenient to express the results in terms of the quantities c_j defined by

$$c_j = \binom{5}{j}^2 \frac{\Delta b_j^2}{\lambda(1-\lambda)} . \quad (6)$$

For our case ($n=5$, $k=8$), one finds that

$$\begin{aligned} c_4 &= 2.4286 \\ c_3 &= 0.8047 \\ c_2 &= 0.1150 \\ c_1 &= 0.0083 \\ c_0 &= 0.0002 \end{aligned} \quad (7)$$

(The quantity b_5 trivially has no spread; it is always equal to unity.) Thus, for example, when $\lambda=1/2$, the spread in b_4 is equal to 0.156, considerably larger than the spread in λ . Each of the b_j 's is distributed according to a binomial distribution, which can be well approximated by a Gaussian distribution.

The next step is to use our knowledge of the distributions of the b_j 's to find the distribution of v . The equation relating the b_j 's to v is the mean field equation Eq. (5).

It is convenient to re-express this equation in terms of the difference δb_j between b_j and its average value $(1-\lambda)$. Thus $b_j = (1-\lambda) + \delta b_j$. In terms of δb_j , Eq. (5) becomes

$$f(v) = \sum_{j=0}^{n-1} g_j(v) \delta b_j, \quad (8)$$

where the functions f and g are defined by

$$f(v) = (v - v^n) - (1 - v^n)(1 - \lambda) \quad (9)$$

$$g_j(v) = v^j (1 - v)^{n-j} \binom{n}{j}. \quad (10)$$

To find the distribution of v , we make one simplifying assumption, namely, that for each set $\{b_j\}$, there is a unique stable solution to Eq. (5). This may not be strictly true, but it appears to be the case for all those sets $\{b_j\}$ which have any appreciable probability.

If this assumption is granted, then one can express the distribution p of v as follows:

$$p(v) = \int p_0(b_0) \dots p_{n-1}(b_{n-1}) \delta(v - v_0(b_0, \dots, b_{n-1})) db_0 \dots db_{n-1}. \quad (11)$$

Here p_j is the distribution of b_j , δ is the Dirac delta function, and v_0 is the solution of Eq. (5). Assuming Gaussian distributions for the b_j 's, one can carry out the integrals in Eq. (11), and the result is an explicit expression for the distribution of v :

$$p(v) = \frac{1}{\sqrt{2\pi}} \left| \frac{d}{dv} \left(\frac{f}{\sigma} \right) \right| \exp \left(- \frac{f^2}{2\sigma^2} \right), \quad v < 1; \quad (12)$$

(probability that $v=1$) = 1 - (probability that $v < 1$).

Here σ is defined by

$$\sigma = \left[\sum_{j=0}^{n-1} g_j^2(v) \Delta b_j^2 \right]^{1/2}. \quad (13)$$

Eqs. (12) and (13) are the main result of this section. They give us the mean-field prediction for the distribution of the density of zeros in terms of known quantities. From this distribution one can easily obtain the distribution of values of the entropy via Eq. (4). Note that the distribution $p(v)$ depends on λ through the function f and through the Δb_j 's.

It is of course very interesting to compare this theoretical distribution with the observed distribution. The graph in Fig. 4 was constructed as follows: For each value of λ , one hundred values of v were chosen at random in accordance with the distribution $p(v)$. Then for each of these values, a corresponding value of the entropy was computed via Eq. (4). The resulting entropy values were plotted. Note that for λ larger than about 1/2, the distribution of entropy values appears to match well the distribution seen in the data (Fig. 1). Even for smaller values the match is not bad, except for one feature: the theoretical prediction does not show the gap just below $H=0.84$ that one sees in the actual data. This disagreement is not surprising. The gap is associated with the transition from periodic to chaotic behavior, and one does not expect the predictions of mean field theory to apply to periodic automata. Only in chaotic automata could one hope to find enough mixing of the symbols for the assumptions of mean field theory to be approximately valid. (See, however, the paper by McIntosh in this volume, which relates automaton behavior not to the prediction of mean field theory but rather to the form of the mean field equation [12].)

5 A Quantitative Comparison

It is satisfying that the theoretical graph looks something like the actual data, but one would also like to have a more quantitative comparison. An obvious way to make such a comparison would be to compute the average entropy and the mean deviation as functions of λ , both for the actual data and for the theoretical distribution. However, the low-entropy periodic automata contribute significantly both to the average and to the deviation, and we do not expect agreement when such automata are involved. Therefore, we would like to be able to eliminate these periodic automata from the average, to see whether the theoretical prediction at least works for the chaotic automata.

On the experimental side, it is straightforward to pick out all the periodic automata and to exclude them from the average. In order to get the corresponding results on the theoretical side, we need to tell the theory how to recognize periodic automata. To do this, we adopt the following cut-off hypothesis: Any automaton for which the mean-field-predicted entropy is less than 0.84 is a periodic automaton, and all others are chaotic. This is a very crude assumption, but it will be interesting to see how well it works. In Fig. 5, we have plotted the theoretical predictions for $\langle H \rangle$ and ΔH , having dropped from the average all automata with a predicted entropy less than 0.84. Also plotted in Fig. 5 are the experimental results with the periodic automata dropped from the average. The agreement is very good. From this agreement we draw two conclusions: First, for chaotic automata, the variation in the mean-field parameters b_j is the principal source of the variation in the entropy. Second, the cut-off hypothesis has some truth to it. (However, it is not strictly true, as we discuss in Section 8.)

6 What happens as n or k gets large

Now that we have good evidence that the source of the variation is in the b_j 's, we can predict how this variation will change as we change the values of n and k . Let us estimate the behavior of the spread in two limits: large n with fixed k , and large k with fixed n . In the latter case, the values of H will typically go as $\log(k)$ --the entropy is greater when there are more symbols--so it makes sense to consider the normalized quantity $H = H/\log(k)$.

The order of magnitude of ΔH is the same as that of the quantity σ defined in Eq. (13). In order to estimate σ , we need to estimate the Δb_j 's for arbitrary n and k . A reasonable estimate is given by $\Delta b_j = \left[\binom{n}{j} (k-1)^{n-j} \right]^{-1/2}$. (The quantity in brackets is the number of ways of filling an n -cell neighborhood, when there are k symbols available and when exactly j of the cells must contain zero.) Using this estimate, and for simplicity letting v equal $1/2$, one finds that

$$\sigma = \left(\frac{1}{2}\right)^n \left[\sum_{j=0}^{n-1} \binom{n}{j} \frac{1}{(k-1)^{n-j}} \right]^{1/2}. \quad (14)$$

If k is held fixed, then this expression decreases exponentially with increasing n . If n is held fixed and k is allowed to get large, the expression is dominated by the term with $j=n-1$, and the whole expression goes as $1/\sqrt{k}$. Thus in both limits, the spread in the entropy goes to zero.

We now have to imagine what happens to the graph of Fig. 1 in each of these limits. Consider first the case in which k gets large while the neighborhood size stays the same. For values of λ to the right of the transition region, the spread in H will presumably get smaller and smaller, in accordance with the above prediction, probably approaching zero as k goes to infinity. Thus the scatter of points will begin to look like a curve. Assuming that this curve remains a curve for all λ , the transition must occur at a definite value of λ if it occurs at all. In the following section we present evidence that there is indeed a sharp phase transition in this limit, occurring around $\lambda=0.27$.

The limit $n \rightarrow \infty$ is likely to be less interesting. As before, the scatter of points will probably become a line, but in this case one may lose the transition altogether. Mean field theory predicts a transition from the quiescent state to a chaotic state at $\lambda=1/n$. (This follows from Eq. (3), with the exponent 5 replaced by n .) This prediction is not correct for small n , but as n increases one expects mean field theory to be more accurate. Thus it is likely that the transition point will be pushed to $\lambda=0$ as n approaches infinity.

7 The Infinite k Limit: An Experimental Test

Let us focus now on the case of a large number of symbols. It is difficult to make a large increase in the size of k experimentally, because the rule table expands as k^5 . Moreover, the spread in H is predicted to go as $1/\sqrt{k}$, as we have just seen, and this is a rather slow convergence. If we used 70 symbols instead of 8, the spread would be reduced only by a factor of 3.

However, it is actually possible to do the experiment when k equals infinity. If there are an infinite number of symbols, then as long as the lattice is finite and the automaton runs for a finite time, there is zero probability that any nonzero symbol will

appear more than once in the entire run. (The symbol zero can appear frequently because one fixes λ at some value less than 1.) The automaton is therefore constantly exploring new parts of the rule table. Moreover, the rule table was constructed at random, so in effect one is simply running a probabilistic automaton with two symbols, zero and "other." An all-zero neighborhood is mapped to zero with probability 1, and any other configuration is mapped to "other" with probability λ .

We have run this binary probabilistic automaton many times with values of λ ranging from 0 to 1/2. (The size and shape of the neighborhood are the same as in our finite- k case.) Fig. 6 shows the resulting values of the entropy. There does appear to be a sharp phase transition at $\lambda=0.27$, and the points fall on a fairly well defined curve, in agreement with our expectations. Note that there does not appear to be a discontinuity in the entropy itself as a function of λ , but rather a discontinuity in the derivative. The existence of this sharp transition, combined with our prediction that ΔH decreases with increasing k , provides strong evidence that the transition region itself shrinks as k increases, the width of the region probably going as $1/\sqrt{k}$.

In collecting the above data, we let the automaton run for 1000 time steps before collecting data. The exact location of the transition changes slightly if one uses a longer waiting time or a larger lattice. (The dependence of a related transition on lattice size has been discussed in reference [13].) Our aim here is not so much to locate the transition point precisely as to show that the transition does become sharp in the infinite k limit. It is interesting to note that as one approaches the transition point from the low λ side, it takes longer and longer for the transients to die out. Thus the transition appears to be of second order.

Sharp phase transitions in probabilistic automata have been observed before. (See for example [14] and [15].) They have also been observed in deterministic inhomogeneous automata with random assignments of rules to sites. (See for example [16], [17], [18], and [13].) Our results show how homogeneous deterministic automata can approach such behavior.

8 A Sharper Transition with $k=8$?

Even if the transition region becomes narrower as k gets larger, there will still be a sizable range of transition points for any realizable finite value of k . If possible, one would like to find a parameter other than λ in terms of which the transition is sharper, even when the number of symbols is held fixed.

Such a parameter is suggested by our cut-off hypothesis. Recall that the hypothesis asserts that an automaton will be periodic or chaotic, depending on whether its mean-field-predicted entropy H_{mf} is smaller or larger than 0.84 bits. Suppose this hypothesis were strictly correct for our automata. Then if one were to plot the actual entropy H against H_{mf} , one would see the following: For values of H_{mf} less than 0.84, the actual entropy would be very small; at $H_{mf}=0.84$, the entropy would suddenly jump to a value close to 0.84 and would rise gradually thereafter. Thus we would have a sharp phase transition. Or instead of using H_{mf} , one could use v_{mf} , the mean-field-predicted density of zeros. Either of these quantiles is computable from the rule table without having to run the automaton.

We therefore propose the following experiment. Let v_{mf} take values from 0 to 1 in steps of, say, 0.01. For each value, choose at random some number of automata whose b_j 's satisfy Eq. (5) with the given value of v_{mf} . Run these automata and compute their entropy. If the cut-off hypothesis were correct, the resulting plot of entropy vs. v_{mf} would show a sharp jump at the value $v_{mf}=0.885$, which corresponds to $H_{mf}=0.84$.

In fact the cut-off hypothesis is not strictly correct. For example, in Fig. 1, one sees a number of low-entropy periodic automata around $\lambda=0.5$ whose existence is not predicted by the theoretical result shown in Fig. 4. The values of H_{mf} for these automata are surely greater than 0.84, and yet the automata have low entropy. However, these automata are relatively rare, so one might still hope to see a reasonably sharp transition in the proposed experiment.

9 Conclusions

Our strongest conclusions are these: As the number of symbols increases, the transition region for our automata almost certainly shrinks, approaching a unique transition point around $\lambda=0.27$ as k goes to infinity. On the other hand, if one increases the size of the neighborhood, the transition region probably gets pushed closer and closer

to $\lambda=0$, so that there is in fact no transition at all in the infinite-neighborhood limit, at least none that can be seen by varying λ .

These conclusions depend on the fact that for the automata we are considering, an all-zero neighborhood always maps to the symbol zero. If this property were removed, it is not clear whether there would be a sharp phase transition in the infinite- k limit. It may also be the case that our results depend crucially on the two-dimensionality of the lattice. To determine how these features of the automata influence the nature of the transition will require more experimental and theoretical work.

Our results also suggest that even for finite k , one can probably sharpen the transition by using a different parameter, v_{mf} instead of λ , to explore the space of automata. This alternative parameter may be useful for finding complex automata near the transition point, which was one of our motivations for looking for a sharper transition. However, there is nothing natural about v_{mf} , so we do not expect any deep insights to emerge through this particular parameterization of the space of automata. In order to see a closer parallel between the cellular automaton transition and physical phase transitions, it may be necessary to use a set of quantities none of which can be computed without actually running the automaton. For example, one might define the "temperature" as the exponential of the average expansion rate. One could then plot entropy against temperature. Both entropy and temperature (by this definition) are measured properties of the automaton's behavior and, unlike λ or v_{mf} , cannot be easily controlled. It would nevertheless be of great theoretical interest, and ultimately of practical value, if one found a sharp phase transition in such a plot.

Another more fundamental question is certainly worth investigating: If there is indeed an analogy between cellular automata and thermodynamics, how many independent thermodynamic variables are there? That is, how many different macroscopic statistical properties of an automaton does one have to specify in order that all other macroscopic statistical properties are then determined? In order to see a sharp phase transition, one would have to be able to vary separately all the independent variables. For example, in order to see a sharp phase transition between liquid water and water vapor, it is not sufficient to plot entropy against temperature while paying no attention to wild variations in the pressure. In a similar way, we may need to pay attention to two or more independent variables in order to see a sharp transition.

Acknowledgments

The ideas presented in this paper have emerged from conversations with a number of people, including David Cai, Doyne Farmer, Howard Gutowitz, Stuart Kauffman, Wentian Li, John Miller, Norman Packard, and Steen Rasmussen. We thank them for their insights and questions. We would also like to thank Stuart Kauffman and Gerard Weisbuch for directing us to some of the work on inhomogeneous automata. Work by WKW was supported in part by the Theoretical Astrophysics group at Los Alamos National Laboratory.

References

1. S. Wolfram, *Physica* 10D (1984), 1.
2. S. Wolfram, *Physica Scripta* T9 (1985), 170.
3. C. G. Langton, *Physica* 22D (1986), 120.
4. C. G. Langton, "Computation at the Edge of Chaos," to appear in *Physica D*.
5. C. G. Langton, Ph. D. thesis, University of Michigan, 1990.
6. W. Li and N. Packard, submitted to *Complex Systems*.
7. W. Li, N. Packard, and C. Langton, in this volume.
8. N. Packard, "Adaptation Toward the Edge of Chaos," Technical Report CCSR-88-5, Center for Complex Systems Research, University of Illinois (1988).
9. S. Wolfram, *Rev. Mod. Phys.* 55 (1983), 601.
10. H. A. Gutowitz, J. D. Victor, and B. W. Knight, *Physica* 28D (1987), 18.
11. Z. Burda, J. Jurkiewicz, and H. Flyvbjerg, preprint NBI-HE-90-04, The Niels Bohr Institute (1990).
12. H. V. McIntosh, in this volume.
13. D. Stauffer, *Physica* 38D (1989), 341.
14. P. Grassberger, F. Krause, and T. von der Twer, *J. Phys. A: Math. Gen.* 17 (1984), L105.
15. W. Kinzel, *Z. Phys.* B58 (1985), 229.

16. B. Derrida and D. Stauffer, *Europhys. Lett.* 2 (1986), 739.
17. G. Weisbuch and D. Stauffer, *J. de Physique* 48 (1987), 11.
18. L. de Arcangelis, *J. Phys. A: Math. Gen.* 20 (1987), L369.

Figure 1

Measured values of single-site entropy for automata with different values of λ . (λ is the fraction of nonzero entries in the rule table.) For each λ , one hundred rule tables were constructed at random. Each automaton was allowed to run for 500 time steps before the entropy was measured.




Figure 2

Density of nonzero symbols as a function of λ , as predicted by simple mean field theory.




Figure 3

The solid curve is the single-site entropy as predicted by simple mean field theory. The dots are the observed values, just as in Fig. 1.




Figure 4

Theoretical distribution of entropy values versus λ . For each λ , one hundred values were chosen at random in accordance with the predicted distribution, Eq. (12). The distribution is similar to the actual distribution shown in Fig. 1.




Figure 5

The solid curve shows the average entropy (upper curve) and the mean deviation of the entropy when periodic automata are dropped from the average. The dashed curves show the theoretical predictions for the same quantities, when all automata with predicted entropies less than 0.84 are dropped from the average.




Figure 6

Entropy versus λ for a probabilistic automaton intended to simulate our deterministic automata in the infinite k limit. The lattice was 128×128 with periodic boundary conditions. Each automaton was allowed to run for 1000 time steps before data was taken. The location of the transition point shifts slightly with larger lattices and longer waiting times.

Fig. 6

$i28 + 160$

