

Emergent Phenomena and Complexity*

Vince Darley
Division of Applied Sciences
Harvard University
33 Oxford Street, Cambridge MA 02138
e-mail: vince@das.harvard.edu

ABSTRACT

I seek to define rigorously the concept of an emergent phenomenon in a complex system, together with its implications for explanation, understanding and prediction in such systems. I argue that in a certain fundamental sense, emergent systems are those in which even perfect knowledge and understanding may give us no predictive information. In them the optimal means of prediction is simulation. I investigate the consequences of this for certain decidability and complexity issues, and then explain why these limitations do not preclude all means of doing interesting science in such systems. I touch upon some recent incorporation of this work into the investigation of self-organised criticalities.

1 MOTIVATION AND OBJECTIVES

The calculations were so elaborate it was very difficult. Now, usually I was the expert at this; I could always tell you what the answer was going to look like, or when I got it I could explain why. But this thing was so complicated I couldn't explain *why* it was like that.

So I told Fermi I was doing this problem, and I started to describe the results. He said, "Wait, before you tell me the result, let me think. It's going to come out like this (he was right), and it's going to come out like this because of so and so. And there's a perfectly obvious explanation for this—"

He was doing what I was supposed to be good at, ten times better. That was quite a lesson to me.

Richard Feynman

The defining characteristic of a complex system is that some of its global behaviours, which are the result of interactions between a large number of relatively simple parts, cannot be predicted simply from the rules of those underlying interactions¹.

The word "simply" in the preceding paragraph is somewhat troublesome. I set out to explain, justify, under-

*Work supported in part by the Kennedy Memorial Trust, London.

¹Chaotic systems (in the mathematical sense of the word) which inherently evade prediction can be incorporated into this definition, by allowing the fact that for such systems we know and understand, in a quantitative way, precisely how and when our predictions will be deficient due to inadequacies in our knowledge of the initial conditions.

stand and rigorise, in the context of such systems, what "simply" means.

With that in mind, let us hypothesise the concept of an 'emergent phenomenon' as a large scale, group behaviour of a system, which doesn't seem to have any clear explanation in terms of the system's constituent parts. Rather than viewing this as the first step in the hierarchical evolution of hyperstructures[1], I am interested in 'first-order emergence' in its own right, as it is currently more amenable to a precise formalisation than the elusive increase in complexity one observes in natural evolution.

Why is it that we seem to be incapable of answering the interesting questions about how and why the final state was reached? In the context of the above quotation, can we always find, or meaningfully postulate, a person like Fermi who can solve our problems through a deeper understanding than our own — side-stepping all of our complicated calculations? More succinctly, do truly emergent phenomena exist?

In order to address these issues, we shall need to clarify what we mean by a 'clear explanation' and what we mean by 'understanding'. This is especially warranted due to the misunderstandings inherent in different people's working definitions of such terms in different fields - complexity and emergence have become very much vogue labels for problems in fields as diverse as economics, artificial life, artificial intelligence, neuroscience, and even in the study of cultural change and development. I will use the term 'emergent system' with the understanding that not all phenomena observed in that system will be emergent.

1.1 Explanation

An explanation is some type of answer to the question 'Why did X happen?'. It can contain any or all of the following characteristics:

- "What precise sequence of events and interactions caused X to happen in this situation?" We argue that we can reduce the problem to its constituent interactions, all of which we can explain (often in a precise, quantitative fashion). Thus we have explained the connection between our initial state and the event X in a causal fashion. This sort of explanation can be used either as pre- or post-explanation of the event².

²Most scientists would not deem this nor its associated mode of understanding as a sufficient characteristic of a definition of either

- “What were the fundamental details of the initial state which caused X ?” This allows us to generalise and determine some characteristics of the class of initial states which bring about X , as compared with the space as a whole. If we can’t generalise, we may reach the conclusion that we were dealing with a special case.

- “Assuming this situation isn’t special in any way, why did X , rather than Y, Z, \dots happen?” This is an elaboration of the previous point. Presumably we can imagine a number of different outcomes to which the system could progress (If we’re so sure it’ll reach X , why aren’t we researching something more interesting). Therefore we would like to have some form of explanation for the mapping between different classes of initial state, and the different outcomes X, Y, Z, \dots

1.2 Understanding

We need to understand the legitimacy of the procedures and applications contained in the previous explanations. Therefore, we either understand, fundamentally, the mathematical tools and techniques used in carrying out the rules of the system or believe in the validity of the body of theory which leads to them. If we are only interested in the abstract concept of emergence then this is a sufficient characterisation.

Otherwise, if we are interested in legitimising the applicability of a particular system to an external field of study, we must understand how that field incorporates the system into its own phenomenology. For instance we may wish to use our results to make statements about physical or economic systems, and in such a case we would like to justify the legitimacy of such claims within an accepted framework of such fields.

Let us now examine successive levels of understanding, in an analogous manner to that for explanation.

- “I understand the rules which govern the actions of every single agent and interaction in the system, precisely.” I could carry out a simulation of the system to explain precisely why (in the first sense above) X happened.

- “I understand the rules and the arena in which they operate sufficiently well that I can make predictions of the outcome very rapidly from the initial state alone, without having to calculate every interaction.” I have some deeper understanding of the system and the legitimising tools, such that I can perform an analysis which reveals some symmetries (probably abstract) which enable me to calculate the outcome more directly. This analysis will presumably reveal at least a partial characterisation of the space of initial states, whose boundaries may have to be sharpened by means of simulation.

- “My analysis and understanding of the system is sufficient to give a clear, precise classification of the space of initial states in terms of the system’s outcome.” Here we have achieved complete success with the previous method of analysis, and developed a mapping from the space of initial states to the space of outcomes described by, for example, a closed-form expression.

‘explanation’ or ‘understanding’.

Now, a simulation of a system which arrives at a given result, and the information contained in the intermediate steps of such a simulation, clearly comprise some form of explanation, which will indeed be very useful to explain some phenomena. If developed in an external fashion, it at least demonstrates sufficient understanding of the original system so as to be able to reproduce its behaviour. This is clearly an important step.

To return to our motivating quotation, however, surely if we *really* understand a system, we shouldn’t need to perform such a simulation. We would understand both the circumstances which are necessary for each relevant phenomenon to arise, and the correlation between the initial state and such circumstances. We could then say, directly, whether and why X, Y or Z will happen.

There is clearly a continuum of levels of difficulty and complexity of analysis. One property of a system will be the minimum level of analysis to which it will yield. To couch this in computational terms, there is a minimum amount of computation which needs to be carried out to predict any given phenomenon. With this in mind, I propose the following definition, which will be made more mathematically precise in the forthcoming sections:

Definition 1 *A true emergent phenomenon is one for which the optimal means of prediction is simulation.*

What we mean by a ‘prediction’ is left deliberately vague - it may be precisely quantitative or just a loose classification. With that proviso, rather than phrasing this definition as a hypothesis about our concept of ‘emergence’, I prefer to use it as a formalisation of the term ‘emergent’ in the discussion which follows. I seek to demonstrate the viability of the above definition as an axiom in the study of emergence in complex adaptive systems. Note that I have yet to demonstrate the existence of emergence.

Emergence as I describe it here has been loosely termed ‘computational emergence’ by [4]. However, abstract computer simulations have been used ever more frequently to gain an understanding of physical phenomena[13, 7] and in the evolution of specific functionality[9]. So computational emergence can, I feel, bear upon Cariani’s ‘thermodynamic emergence’ and ‘emergence relative to a model’.

So, emergent phenomena are those for which the amount of computation necessary for prediction *from an optimal set of rules, classifications and analysis, even derived from an idealised perfect understanding, can never improve upon the amount of computation necessary to simulate the system directly* from our knowledge of the rules of its interactions. I posit that this definition concurs with the wishy-washy heuristic sense people mean when they (ab)use such terms as ‘emergent behaviour’ and ‘emergent complexity’³.

³There are of course some completely un-loaded uses of the word emergence, such as to describe the hardness of a rock, or the colour of an object.

Firstly note that the predictive difficulty of a phenomenon will depend upon both the size, n , and the interaction complexity of a given member of a family of complex systems.

I argue that there is no discontinuous separation between emergence and non-emergence. Emergence is purely the result of a phase change in the amount of computation necessary for optimal prediction of certain phenomena. To explain this phase change, we need to look at the two means of prediction we can use:

- Let $s(n)$ be the amount of computation required to simulate a system, and arrive at a prediction of the given phenomenon. At times we may wish to consider this an idealised quantity referring to the optimally tuned simulation (“God’s simulation”), but always a simulation.

- Our ‘deeper level of understanding’ of the symmetries of the system (in both the obvious and abstract senses), has allowed us to perform a creative analysis and deduce the future state whilst, we hope, circumventing most of the previous-required computation. Let $u(n)$ be the amount of computation required to arrive at the result by this method.

I shall shortly address the issue of defining these ‘amounts of computation’ more precisely, but for the moment let us assume this can be done in a useful manner. Then the above definition states:

$$\begin{aligned} u(n) < s(n) &\Rightarrow \text{the system is non-emergent.} \\ u(n) \geq s(n) &\Rightarrow \text{the system is emergent.} \end{aligned}$$

For simple systems, obviously $u(n) \ll s(n)$ (think of statistical physics), but for many classes of system, as we increase their size and rule complexity, there will be a phase change where the curves $u(n)$ and $s(n)$ cross⁴. The crux is that there is *no discontinuity* separating non-emergent from emergent systems. There is just a phase change in the optimal means of system prediction. Beyond this, perfect understanding of the system does no better than a simulation. All useful predictive knowledge is contained in the accumulation of interactions.

Before proceeding further, let us return to the issue of defining $u(n)$ and $s(n)$. I aim this analysis primarily towards discrete, deterministic systems, although many of the ideas could clearly be made more widely applicable. The obvious problem is that the time required to complete a given computation is machine-dependent. For finite computations the time required by different Turing machines is just related by an arbitrary polynomial. If this phase transition is a real phenomenon, it shouldn’t be machine-dependent. To remove this dependence, we must take into account the complexity of the Turing machine calculating the predictions. I therefore propose the following definition:

Definition 2 *The amount of computation, c , for some process is given by*

$$c = \sum_{s \in \mathcal{S}} (\text{complexity of } s)$$

⁴I shall justify this statement later.

$$= (\# \text{ of steps}) \cdot (\text{complexity of each step})$$

where \mathcal{S} is the set of all steps required in the computation, and the latter equation is true only if all computational steps are the same.

The complexity of a step is measured in terms of the Kolmogorov complexity of a representation of the computation - i.e. the minimal description length.

To justify the machine-independence of this definition, consider a general cellular automaton (which clearly falls within our definition of a complex system), with K possible states for each cell, and a dependence neighbourhood of R cells.

Given $n(R - 1) + 1$ adjacent initial cell-states, we can determine the state of a single cell at the bottom of a light-cone after n time-steps in just $n + (R - 1)n(n - 1)/2 \approx \frac{1}{2}Rn^2$ rule applications. The CA is a mapping from a set of size K^R to one of size K . As this mapping may be completely arbitrary, we can represent each step with some $(\log_2 K^R) \cdot (\log_2 K) = R \cdot (\log_2 K)^2$ bits. So, to leading order, $s(n) = \frac{1}{2}(nR \log_2 K)^2$.

I shall now argue that $s(n)$ is invariant under the following transformations.

1. Pre-calculating a larger table, so that we need simulate only every l ’th step, by using a larger neighbourhood (of size $R + l(R - 1)$).

Clearly $n \rightarrow n/l$, $R \rightarrow R + l(R - 1)$. Thus $s(n)$ is reduced by a factor $(1 - 1/R + 1/l)^2$, which is second order, and only significant in the limit of large memory (l) and small R .

2. Using N machines in parallel to simulate the system. $s(n)$ is obviously unchanged except for a small increase due to the overhead involved in splitting up the problem.
3. Mapping the CA onto another with more states, such that a single cell corresponds to several in the original system, but with smaller neighbourhoods.

Assume we compress a cells into 1, which will now have $K' = K^a$ possible states. $R' \approx R/a$, and $s(n)$ is unchanged.

4. Using *any* Turing machine to simulate the CA.

I shall not attempt to consider a general transformation, but rather appeal to the intuitive idea that increases in Kolmogorov complexity will offset any decrease in computational steps. I postulate that this is true for complex systems in general, not just CA.

Having identified a fundamental interactive unit, and the complexity of the interaction rules for any given system; we can define all computational amounts in the above manner. Let us consider an *idealised* means of prediction in that complex system, requiring an amount of computation $u_{\text{opt}}(n)$, based on perfect understanding.

Definition 3 *The emergence ratio, $\xi = u_{\text{opt}}(n)/s(n)$.*

The emergence coefficient, $\zeta = -\frac{\log \xi}{\log n}$

Where ζ is only really useful if u and s are not exponential in n , but have leading order n^{α_u} and n^{α_s} respectively, in which case $\zeta = \alpha_s - \alpha_u$ to leading order. It is a positive measure of the range of possible levels of predictive understanding we may have in the system ($\zeta = 0 \Rightarrow$ we should simulate; $\zeta > 0 \Rightarrow$ increasingly rapid predictions). A related measure of understanding, which can be useful if we wish to compare how different methods of prediction, $u(n)$, fare as we move around the space of relatively simple systems, is given by:

Definition 4 *Relative understanding*, $\lambda = \frac{u_{\text{opt}}^{(n)} - u(n)}{s(n)}$

Note that, at the phase change, u_{opt} can have a discontinuous gradient. Also, beyond this, $u_{\text{opt}} = s$, and since $u \geq u_{\text{opt}}$, we have that $\lambda \leq 0$. So in an emergent system, our understanding can be at best zero. Our astonishment at the fact that we cannot seem to predict emergent properties, stems not from any failure to understand, but from an inherent property of the system, brought about by the accumulation of interactions. A small reassurance at least.

Now we need no longer deal with any explicit dichotomy between emergent and non-emergent phenomena. The perceived lack of understanding in the former is really just another way of describing the complexity of the map between initial state and final phenomenon. In the sense that a lack of knowledge of the initial conditions will usually cause increasingly poor predictions; this is analogous to a discrete version of chaos.

This explains the apparent paradox of the rules giving a succinct, complete description of the system, whilst still revealing almost nothing about it. Of course any single phenomenon may fall anywhere in the spectrum between trivial prediction and interesting emergence.

Does this necessarily reduce our attempts to understand emergent systems to ‘stamp-collecting’? (In the sense that all we may do is catalogue the results of simulations). Section 4 discusses this issue, demonstrating that there are profitable, interesting possibilities for the study of complex systems.

3 THE ROOTS OF EMERGENCE

Although I have reached the conclusion that we have no useful predictive information, despite a possible ‘perfect’ understanding, I have not yet addressed the issue of the root of emergence and the loss of information. Indeed, can we justify that emergence truly exists?

Many complex systems have been shown to be capable of universal computation (e.g. CA with as few as 14 states). Therefore many questions about their infinite time behaviour are formally undecidable. In the case of finite size or time, what happens to analogues of these undecidable propositions? I posit that they are emergent — the computations retain their irreducibility (see [16, 17, 15] for a heuristic discussion of irreducibility). Therefore the root of emergence and the commensurate loss of information lie in the foundations of decidability and complexity in computation.

Note that these complex systems are often only universal for very particular, rare initial conditions. Following [17], I would suggest that emergence is far more common and occurs in systems which are not computationally universal.

Can we determine, for a given system, whether or not it is emergent? This question seems rather more subtle, now that we have reduced an apparent dichotomy to a continuous parameter range. The answer may be extremely difficult to determine for systems near the transition.

3.1 Decidability of Emergence

Suppose we had an algorithm, which when applied to a complex system could give us a definite answer of ‘Emergent’ or ‘Non-emergent’.

For the case of cellular automata, Wolfram[14] has suggested a classification into four classes of behaviour: homogeneous (Class I), periodic (Class II), chaotic (Class III) and complex (Class IV). Whether a cellular automaton is in class I, II or III has been shown[5, 12] to be undecidable. It follows that the question of emergence in infinite systems is undecidable, with a possible reduction to NP-complete status for finite systems[8].

More heuristically, by considering the halting problem for any complex system which is capable of universal computation, we know that the best (only) means of prediction in such a situation is to ‘run the program’ - i.e. perform the simulation. So all complex systems which fall prey to such isomorphisms would seem to be emergent, and the halting problem then analogises to tell us that, in the general case, the question of whether a system is emergent or not is an undecidable proposition.

The undecidability of emergence presents us with a particularly gnarly problem. For any such undecidable system, we can either:

(i) Assume it is emergent, and use large amounts of computational power for simulations. What do we do if Enrico Fermi suddenly arrives, with his deep understanding, and removes all mystery from the system. We have wasted a huge amount of effort.

(ii) Assume the system is non-emergent, and try to find its deep, hidden inner structure. All this effort could conceivably be in vain.

This devil’s alternative becomes especially important when the system under consideration is one that has been designed by us as a cooperative attempt to solve a particular problem.

4 STUDYING EMERGENCE

I previously defined the emergence ratio ξ . It is a function of the size of the system, n , together with some inherent level of emergence contained in the rules of interaction. It measures the relative efficiency of rule-based ‘understanding’ versus simulation-based ‘understanding’. My results indicate that, in general, we cannot determine this ratio. Some interesting questions arise dealing with the behaviour of ξ with n . In very small systems we can usually perform some form of an analysis, so u will be smaller than s .

As systems become more emergent, and n increases, the propagation of information through accumulated interaction will blur the boundaries of any analysis we try and perform. Trying to generalise in the initial-state space becomes more and more futile, until any previously useful u outstrips s . We gradually approach the worst case – that of being forced simply to classify points in the state-space purely by exhaustive enumeration, with each individual result being determined by a simulation. Generalisation has vanished.

Before moving on to discuss possible approaches to the study of emergent systems, it is important to point out that many systems are not emergent, and therefore amenable to some form of analysis. Such an analysis will certainly not generalise to the emergent complex systems, but is clearly an important and valuable contribution to understanding the context of our investigations – the continuum quantified by the emergence ratio.

For example, it has been demonstrated that in certain highly symmetric classes of one-dimensional cellular automata, the single cell at the bottom of a light-cone after n time-steps can be predicted more quickly than the $n(n-1)/2$ steps of a simulation:

Linear CA[10] have $u \sim n \Rightarrow \xi \simeq 2/n, \zeta \simeq 1$. Quasi-linear CA with radius $1/2$ [11] have $u \sim n^{\frac{\log 3}{\log 2}} \Rightarrow \xi \simeq 2/n^{0.41}, \zeta \simeq 0.41$. The proof of the latter result is particularly informative in the direct manner in which it exploits the symmetry of, for example, the quaternion group.

4.1 Research inside Emergent Systems

I'd like to give an example to suggest that the fact that a given system may lie far beyond the phase change does not mean we should lose all hope in our traditional means of understanding, explanation and prediction:

Consider an analogy with the game of chess. Current computer chess programs use extremely sophisticated but nevertheless brute-force approaches to 'simulate' the game and make predictions. Human grand-masters use a wonderfully subtle combination of pattern-matching, generalisation, comparison and analogy-making with some look-ahead to 'understand' the game and make their predictions. The branching factor of simulations is so high that humans are currently far superior at determining such elusive concepts as positional advantage⁵. I would posit that chess lies on the emergent side of the phase boundary, so that solution by simulation is ultimately the best approach. However, human expertise and understanding seem to achieve an astonishing amount in the face of the above information paradox.

This highlights the fact that my result only states we can do no better than a simulation by using our understanding. A sufficiently sophisticated combination of approaches may be the brain's best bet for prediction in complex systems. The brain itself, of course, is an extremely complex system, which I would suggest lies far beyond the phase change - this clearly has important ram-

⁵Especially if the usual time constraints are removed, as in postal chess.

ifications for artificial intelligence⁶.

Let us now look at the direct study of emergence. There are two alternative view-points from which we could approach emergent complex systems.

Firstly there is the possibility of performing a limited characterisation on the space of emergent systems. Although the boundaries of any such classification in complex system space will be very imprecise, many important systems may fall nicely into such a scheme.

Secondly, given that we know a particular system evolves to a particular type of emergent state under some interesting conditions (or perhaps under most conditions), we may be interested in the behaviour of the system on that emergent state.

1. By virtue of its emergent nature, the interesting, coherent parts of its emergent states will usually be composed of ever changing subsets of the actors which comprise the complex system. Is the split between co-organised and apparently separated actors ergodic? Perhaps we can apply statistical physics techniques to some sub-system.
2. We would like to know the structural stability of the emergent state under external perturbations.
3. The emergent state may exhibit interesting dynamics, e.g. self-organised critical phenomena such as power-law scaling in its structure or constituents.
4. How robust is the emergent state to changes in the underlying rules (interaction stability)?
5. What happens if we allow the system itself to evolve by modification of its own rules?

4.2 Self-Organised Criticality

Let us now look in more detail at just one of the above options⁷:

Definition 5 *A self-organised critical state is a dynamic equilibrium, the frequency of disturbances from which obey a power-law distribution with respect to their size, l :*

$$\begin{aligned} f(l) &= c/l^k \\ \Rightarrow \log f &= c' - k \log l \end{aligned}$$

Where $k > 1$, c is a normalising constant, and these are expectations in the large time limit.

Now, let the set of actors be X ; the parameter under investigation be a_x^t for some $x \in X$ at time t ; and \langle, \rangle be a metric on the parameter space.

Define the average:

$$\bar{a}_X^t = \frac{\sum_{x \in X} \langle 0, a_x^t \rangle}{|X|}$$

⁶The traditional knowledge-based approaches to AI are based completely in conceptual ideas and 'understanding'. My results suggest progress and success by the use of such methods will be harder to achieve than by the use of agent-based and connectionist approaches.

⁷for early discussion on some of these points, see [3, 2].

Then we say a_X is in a critical state if:

$$\mathcal{P}(\lim_{t \rightarrow \infty} \langle a_x^t, \bar{a}_X^t \rangle = l) = cl^{-k}$$

In order to calculate or approximate such a limit for a system, we need to know the rules. In this case all we need is a rule to give us each a_x^t from some local neighbourhood \mathcal{N}_X , at time $t - 1$. We formalise this as follows:

$$a_x^t = \mathcal{R}_{a_x^{t-1}}(\{a_y^{t-1}; y \in \mathcal{N}_X(x)\})$$

Here, in full generality, \mathcal{R} is dependent on all properties of the previous state. For specific systems, the variation in \mathcal{R} may be less general, or even constant.

As an example, let us consider a one dimensional system, with \mathcal{R} constant in time, and a unit neighbourhood.

$$\begin{aligned} a_x^t &= \mathcal{R}_{a_x^{t-1}}(a_{x-1}^{t-1}, a_{x+1}^{t-1}) \\ &= \mathcal{R}'(a_{x-1}^{t-1}, a_x^{t-1}, a_{x+1}^{t-1}) \end{aligned}$$

This is the update equation for a cellular automaton (the derivation of the second line requires that the actors are not uniquely identifiable). Thus CA will fit within this framework, with many other systems, although perhaps very few will satisfy the critical property. This is the kind of phenomenon which is likely to be emergent for many \mathcal{R} .

Current research[6] dealing with one particular type of rule \mathcal{R} — describing the interactions between predictive agents in an artificial economy — has demonstrated the existence of robust self-organised dynamic equilibria. The equilibria are found in the space defined by a metric which isolates the complexity of the predictive algorithms used by the agents. Simulations have shown this is a critical state, with power-law scaling of adaptive changes to the predictive models. Such results will be presented in a forthcoming paper of a less philosophical nature.

5 CONCLUSION

I have given a definition for ‘emergence’, based upon a notion of predictive complexity. Emergence implies computational irreducibility - which can be seen to be the finite analogue of formal undecidability. Hence, the strange lack of understanding in emergent systems has its roots in complexity theory, from which it would seem that emergence itself is an undecidable proposition.

I conclude that not only do emergent systems exist, but also that they match very closely our working definition of the term. Simulation is an optimal means of determining the outcome of such systems, and is thus an important means of investigation. Simulations can be coupled with any of a number of different analysis methods on the emergent state. Together they are suitable for carrying out very interesting, real science in emergent or near-emergent systems, and should lead to explanations of some of the many thought-provoking emergent phenomena we observe. I have elaborated on just one approach,

of current personal interest, to reveal the application of these ideas to critical phenomena.

To summarise, my results suggest the best approach to take in studying emergent complexity is a feedback process of simulation and analysis of the actual emergent phenomena. These results should go some way towards legitimising the concept of a simulation as a real scientific tool for the investigation of *emergence*.

ACKNOWLEDGMENTS

The author would like to thank J. Chalcraft for much insightful and entertaining conversation and criticism.

REFERENCES

- [1] N.A. Baas. Emergence, hierarchies, and hyperstructures. In C. Langton, editor, *Artificial Life III*, pages 515–537, 1993.
- [2] P. Bak, C. Tang, and K. Wiesenfeld. Self-organized criticality. *Physical Review A (General Physics)*, 38(1):364–74, 1988.
- [3] P. Bak, Chao Tang, and K. Wiesenfeld. Self-organized criticality: an explanation of 1/f noise. *Physical Review Letters*, 59(4):381–4, 1987.
- [4] P. Cariani. Emergence and artificial life. In C. Langton, editor, *Artificial Life II*, pages 775–797, 1992.
- [5] K. Culik and S. Yu. Undecidability of CA classification schemes. *Complex Systems*, 2:177, 1988.
- [6] Vincent M. Darley and S. Kauffman. Natural rationality. In W. Brian Arthur, David Lane, and Steven N. Durlauf, editors, *The Economy as an Evolving, Complex System II*. Addison-Wesley, 1997. Also available as Santa Fe Institute working paper 96-08-071.
- [7] J.D. Farmer, S.A. Kauffman, and N.H. Packard. Autocatalytic replication of polymers. *Physica D*, 22:50–67, 1986. *A model of the emergence of autocatalytic networks using the concept of a reaction graph.*
- [8] F. Green. NP-complete problems in cellular automata. *Complex Systems*, 1:453, 1987.
- [9] J. Koza. *Genetic Programming*. MIT Press, 1993.
- [10] O. Martin, A.M. Odlyzko, and S. Wolfram. Algebraic properties of cellular automata. *Commun. Math. Phys.*, 93:219–258, 1984.
- [11] C. Moore. Personal communication, 1993.
- [12] K. Sutner. Classifying circular cellular automata. *Physica D*, 45:386–395, 1990.
- [13] T. Toffoli and N. Margolus. *Cellular Automata Machines*. MIT Press, Cambridge, 1987. *An excellent book on the application of cellular automata to modeling physical systems.*

- [14] S. Wolfram. Stastical mechanics of cellular automata. *Rev. Modern Physics*, 55:601–644, 1983.
Important paper largely responsible for the resurgence of interest in cellular automata.
- [15] S. Wolfram. Origins of randomness in physical systems. *Physical Review Letters*, 55(5), 1985.
- [16] S. Wolfram. Twenty problems in the theory of cellular automata. *Physica Scripta*, T9:170–183, 1985.
- [17] S. Wolfram. Undecidability and intractability in theoretical physics. *Physical Review Letters*, 54(8), 1985.