Artificial Life IV, R. A. Brooks and P. Maes (eds.), pages 394-399, 1994.

Non-Uniform Cellular Automata: Evolution in Rule Space and Formation of Complex Structures

Moshe Sipper
Department of Computer Science
School of Mathematical Sciences
Sackler Faculty of Exact Sciences
Tel Aviv University
Tel Aviv 69978, Israel
e-mail: moshes@math.tau.ac.il

Abstract

Cellular automata are dynamical systems in which space and time are discrete, where each cell obeys the same rule and has a finite number of states. In this paper we study non-uniform cellular automata, i.e. with non-uniform local interaction rules. Two different models are described. In the first a cell's rule may be regarded as a genotype whose phenotypic effect is achieved by rule application. Our focus is on evolution in rule space starting from a random gene pool, i.e. rule population. The second model focuses on the study of complex structures formed by a small number of rules, where the term 'complex' denotes a structure consisting of simple grid cells, acting as a single "organism".

1 Introduction

Cellular automata (CA) are dynamical systems in which space and time are discrete. The states of cells in a regular grid are updated synchronously according to a local interaction rule. Each cell obeys the same rule and has a finite (usually small) number of states (Toffoli and Margolus 1987). The model was originally conceived by John von Neumann in the 1950's (von Neumann 1966).

In this paper we study non-uniform cellular automata, i.e. with non-uniform local interaction rules, an area which seems to have received modest attention (Garzon 1990; Lee et al. 1990; Qian et al. 1990). Our approach is different than these works and is more in the spirit of Artificial Life where cellular automata provide us with "logical universes" (Langton 1986). These are: "synthetic universes defined by simple rules ... One can actually construct them, and watch them evolve." (Toffoli and Margolus 1987). In this context our purpose is to study non-uniform cellular automata with the intent of preserving the three essential features of the original uniform model:

1. Massive parallelism.
2. Locality of cellular interactions.
3. Simplicity of cells (finite state machines).

A major argument in favor of studying non-uniform cellular automata is that due to features (1) and (2), namely massive parallelism and locality of interactions, each cell must retain a copy of the rule in its local memory[1]. Thus we argue that in terms of resources there is no essential difference between uniform and non-uniform automata.

Two slightly different models are described in this paper. In the first a cell's rule may be regarded as a genotype whose phenotypic effect is achieved by rule application. As we shall see a cell's genotype is reproduced if its phenotypic effect promotes fitness. Our focus in this model is on evolution in rule space starting from a *random* gene pool, i.e. rule population.

The second (non-uniform) model introduces a slightly enhanced cellular automaton. We argue that this enhanced automaton is simple enough so that feature (3) (above) is maintained. Our focus here is the study of complex structures formed by a small number of rules. The term 'complex' denotes a structure which consists of simple grid cells, acting as a single "organism".

2 Evolution in rule space

The first model studied is that of binary state, non-uniform cellular automata with a nine cell neighborhood. The initial set-up of each cell's rule table is random where the parameter λ denotes the probability of an entry being one (this is in accordance with the λ parameter introduced by Langton (1986), denoting the percentage of all entries in a rule table which map to non-zero states). Operation of the automaton then proceeds as in the original uniform model, with one difference: evolution takes place not only in state space but also in rule space by having a cell's rule evolve each time the cell is *unsuccessful*, where success may be defined in various ways. Two success criteria discussed ahead are:

1. *Live.* A cell is considered to be successful if it attains a state of one, i.e. "lives". We use the terms alive and dead to represent a state of one and a state of zero, respectively, in accordance with the terminology of the game of Life (Gardner 1970; Berlekamp et al. 1982), one of the best known cellular automata rules.

[1] Although simulations of cellular automata on serial computers may optimize memory requirements by retaining a single copy of the rule this in no way impairs our argument.

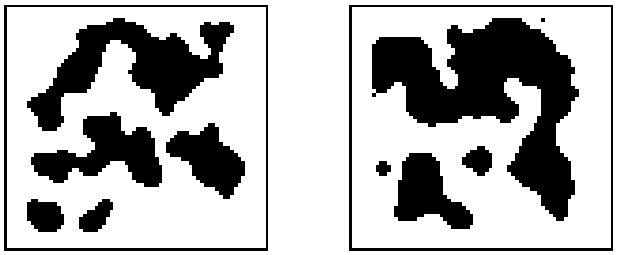

Figure 1: **Territories formed when success criterion *Agree* is employed, with random λ.**

2. *Agree.* A cell is successful if it agrees (i.e. is in the same state) with at least four of its neighbors.

Evolution of an unsuccessful cell's rule is accomplished by selecting one successful neighbor at random and copying its rule (if no successful neighbor exists then the cell's rule remains unchanged). An alternative approach is to copy the most successful rule, a process which is used for non-binary success criteria. Another parameter of the model is whether the copying process is perfect or imperfect,where the latter case is said to involve *mutations*. As noted the initial random population of rules can be viewed as a gene pool, where phenotypic effects are achieved by rule application. In the paragraphs ahead a qualitative presentation of simulation results is given (due to lack of space the actual results are not provided).

We first examine the model using constant λ, i.e. the initial population of rules is generated with the same probability of ones. The success criterion is *Live*. Upon studying the simulation results we observe that for all values of λ the rule grid converges[2]. Furthermore, two interesting thresholds emerge, namely $\lambda = 0.8$ and $\lambda = 0.6$. The first value may be termed the "threshold of life", above which grid cells are guaranteed to attain a state of one, i.e. live. Life is attained not by one rule (which is possible only for $\lambda = 1$) but by a *coalition* of rules, demonstrating an *emergent* behavior of the model. The second threshold is that of $\lambda = 0.6$, and may be termed the "coalition threshold", above which coalitions of rules are formed, while below it one rule emerges as the winner. We also experimented with a population of rules generated with some constant probability λ with a small number of rules generated with a higher λ. It was observed that even a small number of higher λ cell rules is sufficient to induce proliferation of the entire rule grid.

The next case studied is one in which each cell rule is created by first generating a random λ with uniform distribution in the range [0..1]. This λ is then used to generate the cell's (random) rule. Here convergence of the rule grid is much more rapid than the constant λ case and is actually logarithmic. When success criterion *Agree* is employed (with random λ) the surviving rules form "territories" of live cells (see Figure 1).

What happens when the rule copying process is imperfect, i.e. when mutations are involved? In a population of rules with random λ convergence is extremely rapid and the final configuration is one where all cells are alive (success criterion is *Live*). When constant λ is used convergence is much slower though the final configuration is again one in which all cells are alive (unless the mutation probability is too small).

The above criteria of success, namely *Live* and *Agree* admit many local minima in rule space which are all equally valid as far as the evolutionary process is concerned. We noted that coalitions of rules are formed which conform to one of these minima.

It is natural to inquire as to what happens when the success criterion is such that a global minimum exists. We examined one such criterion, namely *Parity*, where a cell is successful if it is equal to the parity of its neighbors in the previous time step (the parity of a cell is equal to 0 if it has an even number of live neighbors, 1 otherwise) (Toffoli and Margolus 1987). Our results indicated that a global minimum is indeed reached.

As a final example we consider a success criterion which is non-binary. This is the Iterated Prisoner's Dilemma (IPD) discussed extensively by Axelrod (1984). Each cell plays IPD with its neighbors where a value of one represents cooperation and a value of zero represents defection. In this case a cell copies (with a small probability of mutation) the rule of the neighboring cell with the highest ranking total payoff (computed by summing the eight individual payoffs). In the trial runs of IPD convergence to a single rule occurred , however each time to a different one. Thus, it is evident that this criterion admits many local minima. An interesting phenomenon becomes apparent upon examining the winning rules: the average percentage of ones is 60%, i.e. cooperation is preferred. This value is close to that of the successful TIT-FOR-TAT strategy (Axelrod 1984) whose percentage of ones is 64%[3].

3 Formation of complex structures

The non-uniform automaton model considered in this section is an enhancement of the first model. Each cell is either *vacant*, containing no rule, or *operational* consisting of a finite state automaton which can, in one time step:

1. Access its own state and that of its immediate neighbors.

2. Change its state, or the state of an immediate neighbor. If a cell's state is changed by more than one cell, contention occurs which may be resolved either randomly or deterministically (i.e. defined by the rules). The rules presented in this paper do not admit such contention.

3. Copy its rule onto a neighoring vacant cell. A special case is cell rule mobility where one copy is made in an

[2]The term *rule grid* denotes the grid of cell rules whereas the term *grid* denotes the grid of cell states.

[3]A value computed by assuming that ties (i.e. an equal number of cooperating neighbors and defecting ones) are broken in favor of cooperation. This choice is based on one of the qualities of a "good" strategy discussed by Axelrod (1984), namely the quality of *forgiveness*.

adjoining cell and the cell's own rule is erased (i.e. the cell becomes vacant). Again contention (in this case the copying of more than one rule onto the same cell) may be resolved either randomly or deterministically. In this paper this type of contention occurs when two cells attempt to move to the same cell, and this is resolved randomly (i.e. one cell "wins" while the other moves to a different vacant cell).

4. A cell may contain a *small* number of different rules. At a given moment only one rule is *active* and determines the cell's function. A non-active rule may be activated or copied onto a neighoring cell.

Our main additions to the original model, aside from non-uniformity, are in allowing an automaton to change the state of its neighboring cells and to copy itself onto them. In the rest of this paper we consider automata with a nine cell neighborhood and three possible grid states, denoted $\{0, 1, b\}$. Note that a vacant cell may be in any grid state as it can be changed by operational neighboring cells. Our focus in this section is the study of complex structures formed by a small number of rules (due to lack of space the rules are not provided).

3.1 A self-reproducing loop

Our first example involves a simple self-reproducing loop motivated by Langton's work (1986,1984) who described such a structure in uniform cellular automata. Langton's loop (motivated by Codd (1968)) makes dual use of the information contained in a description to reproduce itself. The structure consists of a looped pathway, containing instructions, with a construction arm projecting out from it. Upon encountering the arm junction the instruction is replicated, with one copy propagating back around the loop again and the other copy propagating down the construction arm, where it is translated as an instruction when it reaches the end of the arm.

The important issue to note is the two different uses of information, interpreted and uninterpreted, which also occur in natural self-reproduction, the former being the process of *translation*, and the latter *transcription*. In Langton's loops translation is accomplished when the instruction signals are "executed" as they reach the end of the construction arms, and upon the collision of signals with other signals. Transcription is accomplished by the duplication of signals at the arm junctions (Langton 1984).

The loop considered in this section consists of five cells and reproduces within six time steps . The initial configuration consists of a grid of vacant cells (i.e. with no rule) with a single loop composed of five cells in state 1, each containing the (same) loop rule (Figure 2a). The arm extends itself by copying its rule onto an adjoining cell, coupled with a state change to that cell. The new configuration then acts as data to the arm, thereby providing the description by which the loop form is replicated. When a loop finds itself blocked by other loops it "dies" by retracting the construction arm. Figure 2b shows the configuration after several time steps.

In his paper Langton (1984) compares the self-reproducing loop with the works of von Neumann (1966)

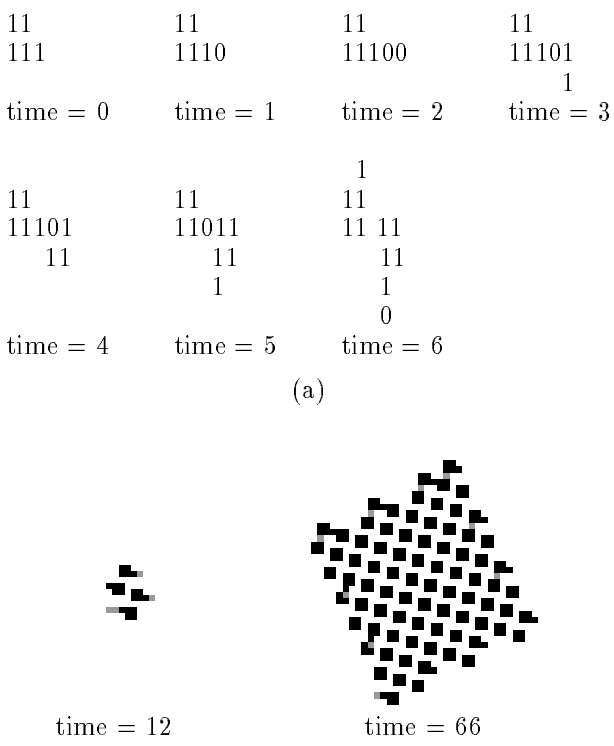

Figure 2: **Self reproducing loop.**

and Codd (1968), drawing the conclusion that although the capacity for universal construction, presented by both, is a *sufficient* condition for self-reproduction it is not a *necessary* one. Furthermore, as Langton points out, naturally self-reproducing systems are not capable of universal construction. His intent was therefore to present a simpler system that exhibits non-trivial self-reproduction. This was accomplished by constructing a rule in an eight-state cellular automaton, in which the dual nature of information, i.e. translation and transcription is utilized.

In the loop presented above simple transcription is accomplished as an integral part of a cell's operation, since a rule can be copied, i.e. treated as data. Once a rule is activated it begins to function by changing states in accordance with the grid configuration, thereby performing translation on the surrounding cells (data). Essentially, the loop operates by transcribing itself onto a neighboring cell while simultaneously writing instructions (in the form of grid states) that will be carried out at the next time step.

In Langton's system each grid cell initially contains the rule that supports replication whereas in our case the grid cells are vacant and the loop itself contains all the information needed. In both cases reproduction is not

coded entirely into the "transition physics" but rather is "actively directed by the configuration itself" where "the structure may take advantage of certain properties of the transition function physics of the cellular space" (Langton 1984). Thus interest in such systems arises since they display an interplay of active structures taking advantage of the characteristics of cellular space.

3.2 Reproduction by copier cells

In the previous section we described a self-reproducing loop, which exhibited a two-fold utilization of information, i.e. translation and transcription. In this section we examine a model of reproduction consisting of passive structures copied by active (mobile) cells. The motivation for our approach lies in the information flow in protein synthesis, where passive mRNA structures are translated into amino acids by active tRNA cells. Each tRNA cell matches one specific codon in the mRNA structure and synthesizes one amino acid.

Our system consists of stationary structures composed of vacant grid cells comprising the passive data to be copied. The copy ("synthesis") process is accomplished by three types of copier cells, denoted X, Y, and Z which are mobile units, "swimming" on the grid, seeking an appropriate match (remember that cellular mobility is possible by using rule copying). When such a match occurs the cell proceeds to create the appropriate substructure, as in the case of a tRNA cell synthesizing the appropriate amino acid. The final result is a copy of the original structure.

The process is demonstrated in Figure 3. The initial configuration consists of a passive structure coupled with X,Y and Z cells randomly distributed on the grid (Figure 3, $time = 0$). Each time step these copier cells move to a neighboring vacant cell (shown as blank squares) at random, unless a match is found which triggers the synthesis process. Each of the three copiers matches exactly one codon, which is a structure composed of three (passive) cells. Figure 3 ($time = 435$) shows the process at an intermediate stage and at the final stage ($time = 813$) where the copy has been produced.

The copy created is not an exact duplicate but rather a "complementary" one. The reason for this is that we wish to avoid endless copying which would occur had an exact duplicate been created. Since our model is inherently local we cannot maintain a global variable specifying that the synthesis process has been completed. The only way to avoid an endless chain of duplicate substructures is by locally specifying that a copy has been completed. This is accomplished by creating a complementary sub-structure, which does not match any copier cell and is not duplicated further.

3.3 Formation and replication of complex organisms

The final system presented involves the formation and replication of complex structures which are created from grid cells and behave as single "organisms" once formed. The system consists initially of two cell types, builders (A cells) and replicators (B cells), floating around on the grid.

Figure 4 demonstrates the operation of the system. At time 0 A and B cells are distributed randomly on the grid and there are two vacant cells in state 1 acting as the core of the building process. The A cells act as builders by attaching ones at both ends of the growing structure. Once a B cell attaches itself growth stops at that end (time 111).

When a B cell attaches itself to the upper end of a structure already possessing one zero a C cell is *spawned*, which travels down the length of the structure to the other end. If that end is as yet uncompleted the C cell simply waits for its completion (time 172). The C cell then moves up the structure, duplicating its right half which is also moved one cell to the right (time 179). Once the C cell reaches the upper end it travels down the structure, spawns a D cell at the bottom and begins traveling upward while duplicating and moving the right half (time 187). Meanwhile the D cell travels upwards between two halves of the structure and joins them together (time 190).

This process is then repeated. The C cell travels up and down the right side of the structure, creating a duplicate half on its way up. As it reaches the bottom end a D cell is spawned which travels upward between two disjoint halves and joins them together. Since joining two halves occurs every second pass the D cell dies immediately every other pass (e.g. time 195).

There are interesting features to be noted in the process presented. Replication should begin only after the organism is completely formed. However there can be no global indicator that such a situation has occurred (see also Section 3.2). Our solution is therefore local: a B cell upon encountering an upper end which already has one zero completes the formation of that end and releases a C cell which travels down the length of the structure. This cell will seek the bottom end or *wait* for its completion. Only at such time when the structure is complete will the C cell begin the replication process.

Replication involves two cells operating in unison where the C cell duplicates half of the structure while the D cell "glues" two halves together. Again it is crucial that the whole process be local in nature since no global indicators can be used.

The spawning of C and D cells are provided for by our model since as noted above a cell may contain a small number of different rules, where only one is active at a given moment. Therefore, the initial B cells can contain all three rules: B,C,D.

The design of our system is even more efficient than that however, requiring only two rule tables, one for A cells and one for $B/C/D$ cells. Each entry of the $B/C/D$ rule table is only used by one of the cell types (i.e. the entries are mutually exclusive). At a given moment the cell has one active rule (which determines its type). If the table entry to be accessed belongs to the active rule-it is used, otherwise a default state change occurs. The default transformation is a move to a random vacant cell for B cells and no change for C and D cells. This may

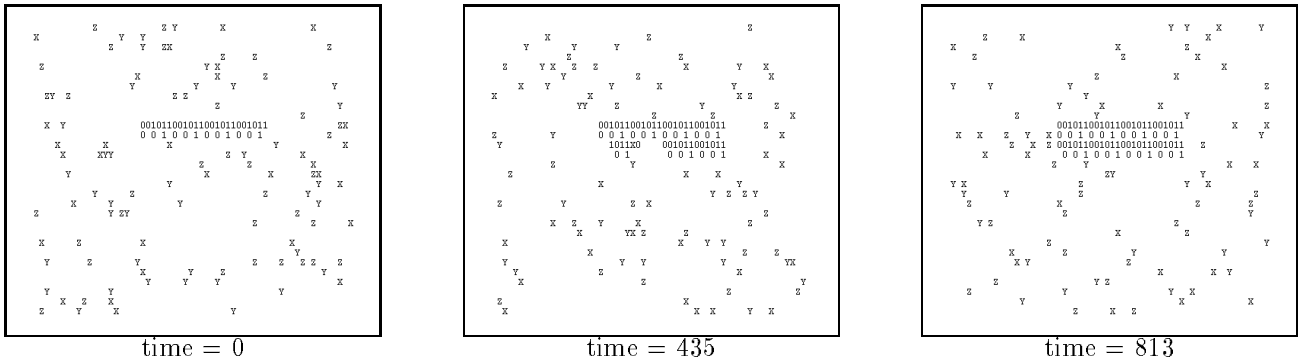

time = 0 time = 435 time = 813

Figure 3: **Reproduction by copier cells.**

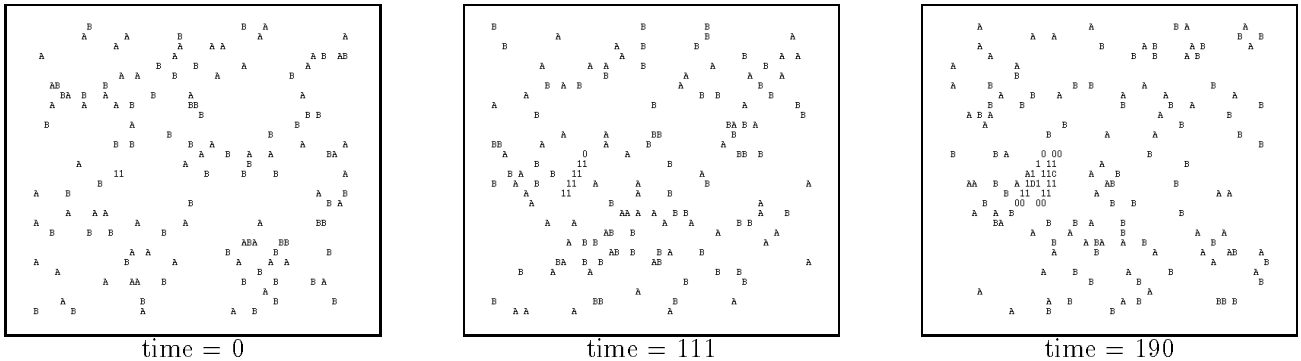

time = 0 time = 111 time = 190

time = 195 time = 323 time = 342

(a) Overview of process.

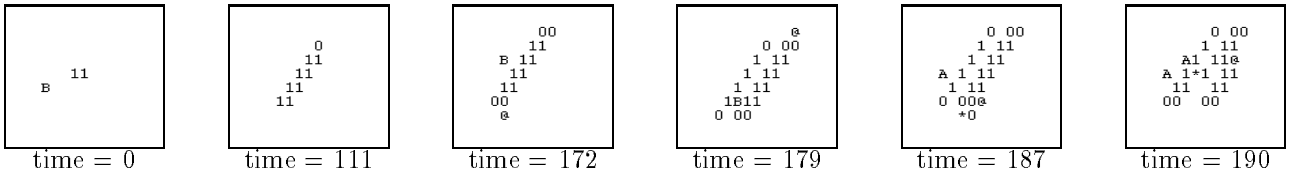

time = 0 time = 111 time = 172 time = 179 time = 187 time = 190

(b) Zoom of intermediate stages. C cells are represented by @, D cells by $*$.

Figure 4: **Formation and reproduction of complex organisms.**

be regarded as a form of differentiation where the cell contains the entire rule table (DNA) but uses only those parts which are relevant to its current functioning.

4 Discussion

The initial random population of rules in the first model (Section 2) can be viewed as a gene pool, where phenotypic effects are achieved by rule application. We observed how successful (fit) genes proliferate the population, in some cases forming coalitions, in others one gene emerges as the winner.

Observing our results we note that when the initial population of rules is generated with random λ convergence is much more rapid than with constant λ. An interesting analogy may be drawn with the biological phenomenon of sex, for which no accepted theory exists. One hypothesis suggests that animal sexuality helps diversify the gene pool thereby promoting more rapid adaptation (Hamilton et al. 1990). In our case diversification is achieved by using random λ which indeed promotes rapid convergence (adaptation in our model).

In his book Dawkins (1986) discusses the issue of possessing "5% of an eye" in relation with objections to the theory of evolution. He argues convincingly that 5% vision is better than no vision at all. Some of our experiments involved a small number of cells which are 5% "more fit" than the rest. We observed that in a majority of the trials such a gene (rule) came to dominate the gene pool. Although our model is simplified in relation to real life it nonetheless demonstrates how even a small advantage is crucial in the "survival" race.

The formation of territories when the success criterion is *Agree* can be regarded as a simple form of epistasis[4]. A gene's success depends on the interaction with its neighboring genes, as opposed to , say, the *Live* criterion where a cell's fitness depends solely on its own gene.

In the second model (Section 3) we concentrated on the formation of complex organisms, composed of simple grid cells. These are formed by using only a small number of rules. The examples presented demonstrate a main feature of our work, namely the power it offers in creating models of interest. As opposed to uniform cellular automata, where each cell contains the same automaton our model enables the creation of systems in which there are different automatons operating in unison. This is especially noted in our last example where complex structures were formed by cooperative operation. Our model also enables the simplification of processes which in uniform automata require complex rule tables with several states per cell. This is noted upon examining the self-reproducing loop which consists of only five cells yet promotes its own replication.

The model presented seems to offer fertile grounds for further investigation. Complex structures may be formed, exhibiting real life properties. Indeed, the dynamic behavior of the last two systems is somewhat reminiscent of observations of organisms under a microscope. As discussed by Langton (1986) such complex structures are essentially virtual state machines, i.e. higher order automata composed of lower order ones, where first order automata are those that occupy every cell and serve as the basic building blocks.

The work presented in this paper suggests a practical model for studying Artificial Life phenomena consisting of enhanced non-uniform cellular automata which evolve in rule space as well as state space. The issues to be explored involve the evolution of complex structures, where the diversity offered by our model coupled with its simplicity seem to present us with a viable system for such explorations.

[4] In the context of Artificial Life this means any interaction between genes, i.e. the extent to which the contribution to fitness of one gene depends on the values of other genes.

References

Axelrod, R. (1984). *The Evolution of Cooperation.* New-York: Basic Books, Inc.

Berlekamp, E. R., J. H. Conway, and R. K. Guy (1982). *Winning ways for your mathematical plays*, Volume 2, Chapter 25, pp. 817–850. New York: Academic Press.

Codd, E. F. (1968). *Cellular Automata.* New York: Academic Press.

Dawkins, R. (1986). *The Blind Watchmaker.* W.W. Norton and Company.

Gardner, M. (1970, October). The fantastic combinations of John Conway's new solitaire game "life". *Scientific American*, 120–123.

Garzon, M. (1990). Cellular automata and discrete neural networks. *Physica D 45*, 431–440.

Hamilton, W. D., R. Axelrod, and R. Tanese (1990, May). Sexual reproduction as an adaptation to resist parasites: a review. *Proceedings of the National Academy of Sciences USA 87*, 3566–3573.

Langton, C. G. (1984). Self-reproduction in cellular automata. *Physica D 10*, 135–144.

Langton, C. G. (1986). Studying artificial life with cellular automata. *Physica D 22*, 120–140.

Lee, Y. C., S. Qian, R. D. Jones, C. W. Barnes, G. W. Flake, M. K. O'Rourke, K. Lee, H. H. Chen, G. Z. Sun, Y. Q. Zhang, D. Chen, and C. L. Giles (1990). Adaptive stochastic cellular automata: theory. *Physica D 45*, 159–180.

Qian, S., Y. C. Lee, R. D. Jones, C. W. Barnes, G. W. Flake, M. K. O'Rourke, K. Lee, H. H. Chen, G. Z. Sun, Y. Q. Zhang, D. Chen, and C. L. Giles (1990). Adaptive stochastic cellular automata: applications. *Physica D 45*, 181–188.

Toffoli, T. and N. Margolus (1987). *Cellular Automata Machines.* Cambridge, Massachusetts: The MIT Press.

von Neumann, J. (1966). *The Theory of Self-Reproducing Automata.* Illinois: University of Illinois Press. Edited and completed by A.W. Burks.