

Redrawing the Boundary between Organism and Environment

Tim Taylor

Institute of Perception, Action and Behaviour, University of Edinburgh
JCMB, The King's Buildings, Mayfield Road, Edinburgh EH9 3JZ, U.K.
tim.taylor@ed.ac.uk

Supplementary material available at <http://homepages.inf.ed.ac.uk/timt/papers/rboe/>

Abstract

In this position paper, I argue that a fruitful, and as yet largely unexplored, avenue for artificial life research lies in modelling organisms (specifically, phenotypes) and environment as a single dynamical system. From this perspective, the origin and evolution of life is the progressive control of the dynamical system at a local level by constraints which are represented on an organism's genome. Such an approach shifts the focus of artificial life models away from the design of individuals, towards the *interaction* of an individual with its dynamic environment. It also blurs the boundary between organism and environment; the most important modelling distinction is no longer between an organism's body and its external environment, but rather between the genome (which is treated as an essentially symbolic structure) and phenotype-plus-environment combined. An evolutionary cellular automata system, called EvoCA, is introduced as a tool to explore these ideas. To demonstrate how this approach differs from traditional studies, two example applications of EvoCA are presented. One concerns sensor and effector evolution; the other concerns the origin of life, and in particular the evolution of genome-regulated self-stabilising dynamics. Advantages of the new approach are summarised, and some potential criticisms are considered. The paper concludes with a discussion of some implications of this shift in perspective.

Introduction

One of the key challenges facing artificial life researchers, as well as biologists, is to explain the origin of living organisms from a non-living environment (Bedau et al., 2000; Maynard Smith, 1986). Furthermore, in order to build artificial evolutionary systems, we would like to know how to produce highly evolvable systems, in which agents can control and exploit their environment in unlimited and increasingly complex ways.

Most ALife work on the evolution of life has employed a strong representational distinction between living organisms and their environment. Examples include *Tierra* (Ray, 1991) and *PolyWorld* (Yaeger, 1994). In *Tierra*, for instance, individuals are computer programs with associated instruction pointers, registers, stacks, etc. Interactions between an individual and its environment can only be achieved in a limited number of predefined ways, such as by the allocation of

memory in order to reproduce (an interaction with the abiotic environment), or by reading machine instructions from a neighbouring program (an interaction with the biotic environment). In these systems the environment is often modelled as a rather inert medium, the only significant role of which is just to provide a "place" in which organisms exist. For further discussion of this topic, see (Taylor, 2001).

Even in work where no such distinction exists between organisms and environment, individuals, and the dynamical laws of the environment, are carefully crafted to achieve a particular type of behaviour. Examples of this type include von Neumann's self-reproducing automata (von Neumann, 1966), simulations of autopoietic systems (Varela et al., 1974; McMullin and Varela, 1997), and Holland's α -Universes (Holland, 1976).

Neither of these approaches—using a strong representational distinction between living and non-living entities, or carefully crafting the "laws of physics" of the world for a particular purpose—can provide much insight of how life first originated from a non-living environment which, presumably, was not specifically designed to support it.

Much of this work is characterised by an emphasis on the computational capacities of the organisms. The *Tierran* language, for example, is computationally universal, but this does not mean that *Tierran* programs can interact with their environment in unlimited ways. Accompanying this perspective has been an (over-)emphasis on the process of self-reproduction, often to the exclusion of other important issues, such as the properties of the environment, and the representational relationship between organisms and environment. If nothing else, the poor evolvability of these systems demonstrates that the processes of self-reproduction with heritable mutation and selection, by themselves, are insufficient to explain the evolutionary origin of complexity.

Howard Pattee, a physicist by training, has devoted much of his career to the question of the origin of life (Rocha, 2001). His particular perspective is the issue of how semiotics (i.e. symbol systems, such as genomes, and their associated semantics in the context of an organism) can originate from a purely physical environment.

Copyrighted Material

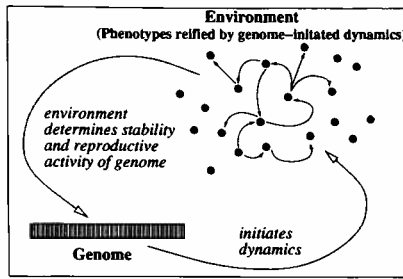


Figure 1: Semantic Closure: Closing the loop between genotype, phenotype and environment

Pattee argues that the distinction between the material and symbolic aspects of living organisms, seen as an example of the more general epistemological distinction between laws and initial conditions, is a defining feature of life, and also a necessary condition for open-ended evolution (Pattee, 1995a; Pattee, 1995b). He explains the relationship between the two as follows:

Writing symbols is a time-dependent dynamic activity that leaves time-independent structure or record. ... Symbols are read when these structures re-enter the dynamics of laws as constraints. Any highly evolved formal symbol system may be viewed as a particularly versatile collection of initial conditions or constraints, often stored in a memory, producing significant or functional behavior that is usefully described by locally selected rules rather than physical laws. ... [A]ll symbol systems must have material embodiments that obey physical laws. But for the reasons just stated, the lawful description of symbols, even though correct in all details, can reveal no significance. (Pattee, 1995b)

The symbols recorded on the genome ultimately acquire semantics in an organism in the context of the survival value of the dynamics that they initiate (i.e. natural selection of phenotypes). It is this autonomous structure-function self-referent organisation that is entailed in Pattee's term "semantic closure" (Figure 1).

This perspective, then, sees organisms as entities whose phenotypes are embedded within an environment viewed as a dynamical system, and whose genotypes interact with the environment by specifying constraints¹ upon its dynamics, thereby generating the phenotypes. That is, the abiotic environment has its own dynamics and self-organisational properties; genotypes act to "sculpt" these pre-existing dynamics by supplying constraints. From this point of view, the most important distinction is not between organisms and their abiotic environment, but rather between the environment as a

¹Throughout this paper the general term 'constraint' is used to cover initial conditions, constraints and boundary conditions. For further discussion of these concepts, see (Pattee, 1995a)




Figure 2: Schematic illustration of EvoCA's design. This specific example shows EvoCA-B. See text for details.

whole (including organism phenotypes) and organism genotypes. It is the relatively time-independent genotypes, by supplying local constraints to the dynamics of the environment, that reify phenotypes as distinct entities within the environment.

The EvoCA System

A simulation platform, called EvoCA, was designed to explore this dynamical systems view of organisms and environments.² EvoCA is built upon a cellular automaton (CA) system. Cellular automata were chosen because they are fairly simple, discrete time and space dynamical systems, whose behaviour has been extensively studied.

The system consists of two layers (see Figure 2). Layer 1 is the environment, modelled as a standard CA. The system can be adapted to use any kind of CA; the examples described here use a 2D environment, with two possible states per cell, and the standard Game of Life update rules. Layer 2 is a discrete grid of the same dimensions as Layer 1. Each grid position in Layer 2 can contain zero or one genomes, which will be described later. The action of a genome is centred upon the corresponding grid position in Layer 1. Layer 2 has no dynamics as such (i.e. genomes are relatively inert structures).

An evolutionary algorithm is used to evolve populations of genomes, with selection based upon the dynamics that a genome produces.³ Two distinct versions of EvoCA exist, which employ different kinds of evolutionary algorithm: EvoCA-A (abiotic selection) uses a standard generational genetic algorithm, while EvoCA-B (biotic selection) uses a natural selection algorithm. The details of genome design and action are also different in each system. The two systems are described in the following sections, and typical results from each are presented in order to demonstrate particular aspects of their behaviour. As these results are included

²The source code is included in the supplementary material.

³Various authors have experimented with evolving the transition function of a CA to achieve a particular task, e.g. (Crutchfield and Mitchell, 1995). This is fundamentally different to the current approach as it entails evolving the "laws of physics" of the environment rather than constraints to control a given set of laws.

solely for the purposes of illustrating certain features and consequences of the general modelling approach being advocated, full details of the experiments are omitted. These details are included in the supplementary materials.

EvoCA-A

Genomes A genome in EvoCA-A comprises a variable length list of genes. Two types of gene are available: timed and conditional. Both types specify a particular target cell in Layer 1 (whose position is defined relative to that of the genome) and a target state for that cell. A maximum radius is defined for each dimension of the CA to confine the position of the target cell relative to the genome.

Each gene additionally specifies a precondition that must be satisfied in order for it to be activated. Timed and conditional genes have different types of preconditions. Timed genes specify a time (i.e. a specific iteration of the CA) at which they act. At the specified iteration, the gene sets the state of the target cell to the target state. Conditional genes specify a watch cell and watch state. The watch cell specification is confined to the set of cells that are direct neighbours of the target cell. Whenever the specified watch cell is in the specified watch state, the conditional gene is triggered, setting the state of the target cell to the target state.

Every gene in the genome is checked at each iteration of the CA to see whether it should be activated for that iteration. Whenever any gene is activated, its action overrides the normal CA transition function for the target cell for that particular iteration.

The Evolutionary Algorithm A standard generational genetic algorithm is used to evolve a population of individuals. Each individual is evaluated in isolation, and placed in the same grid position in Layer 2. The iteration count of the CA is reset to zero at the start of each evaluation. All cells are initially set to the quiescent state, except those which have non-quiescent states specified by timed genes acting at time zero, or those that are influenced by external signals (described later). The CA is then allowed to run for a given number of iterations, with the genes of the genome setting specific cell states when they become active.

The fitness function depends upon the particular task in hand; an example is given in the next section. In addition to one-point crossover and gene mutation, a number of other genetic operators are also available: gene insertion (a random gene is inserted into an existing genome); gene deletion (an existing gene is deleted from a genome); gene reversal (the order of a sequence of genes between two selected points in the genome is reversed); and gene duplication (a sequence of genes between two selected points in the genome is duplicated at the end of the genome). A limit on the maximum allowable genome length is defined.

Example Application: Sensor and Effector Evolution

The following example demonstrates the application of the




Figure 3: Goal pattern for EvoCA-A experiments

approach to the topic of sensor and effector evolution. The system was set up as shown in Figure 3, using a 2D toroidal CA of size 75x75 cells. To evaluate a genome, it was placed in the centre of the genome area shown. The maximum radius of gene action is shown by the boundary of this area (i.e. genes could directly set the state of only those cells within the genome area). In addition, two goal areas were defined, along with two signal positions, as shown. Note that the signal positions adjoin, but do not overlap, the genome area.

The task that the organisms were set is as follows. Each genome was tested under three conditions: left signal, right signal, and no signal. For the left signal condition, the cells in the 2x2 left signal position (Figure 3) were initially set to state 1 (the non-quiescent state). In this condition, the task of the organism was to produce activity in the left goal area. Any activity in this goal area over the course of the evaluation was rewarded, with the darker shaded cells in Figure 3 being rewarded more than the lighter cells in the goal area. However, any activity in the right goal area was penalised.

For the right signal condition, the opposite signal positions and goal areas were used in the fitness calculation. For the no signal condition, any activity in either goal area was penalised. The scores from the three conditions were combined to produce a final fitness value for the organism. Full details are given in the supplementary materials.

Results In all experiments the system was able to evolve organisms that could perform the task well (i.e. they initiated activity in the appropriate goal area, and no activity in the other goal). In some runs, the best evolved individual would produce a “glider” (a moving pattern of activation) that travelled from the genome area to the appropriate goal area. In other runs, more extensive and complex patterns of activation were observed; an example is shown in Figure 4. The implications of these results are discussed later.

EvoCA-B

EvoCA-B has modifications to allow natural (biotic) selection of organisms, rather than the artificial selection method (the genetic algorithm) used in EvoCA-A. In EvoCA-B, the whole population of genomes exists in Layer 2 concurrently, and the survival and reproduction of each genome is deter-




Figure 4: EvoCA-A: Sensor and effector evolution. Snapshots of the activation of the CA towards the end of the evaluation period for each of the three signal conditions: (a) Left signal (b) No signal (c) Right signal. For reference, the goal areas, genome area and signal areas are shown in gray.

mined by its local environment in Layer 1 (which can be influenced by the genome—and by other nearby genomes). A schematic illustration of the design is shown in Figure 2.

Genomes For practical implementational reasons, genomes in EvoCA-B are somewhat simplified. A genome can only set the state of a *single* target cell in Layer 1 corresponding to the genome’s position in Layer 2, rather than being able to set states of target cells over an extended area as is the case in EvoCA-A. The genome consists of conditional genes only (no timed genes are used). The condition section of a gene specifies a combination of the current state of the associated cell, plus the *number* of neighbouring cells that are currently in a non-quiescent state. So, for 2 states and 8 neighbours, there are $2 \times (8+1) = 18$ possible conditions. If a gene’s condition is met, then the gene sets the state of the cell as before. In addition, a gene also encodes a direction for the genome to move on the Layer 2 grid if the gene’s condition is met (which could be to one of the directly neighbouring cells, or to stay put). Thus the genomes are acting in a manner reminiscent of Turing machines, reading and writing “data” from Layer 1 and conditionally moving position at each iteration.⁴

The Evolutionary Algorithm Each possible local configuration (i.e. the configuration of the states of a cell and its 8 immediate neighbours) has a genome reproduction probability associated with it. Similarly, each local configuration has associated with it a genome death probability and a genome reproduction-fidelity probability. (All these probabilities are hard-wired and constant.) For a genome at a given moment in time, its probability of reproduction (and, if it does reproduce, its probability of producing a faithful copy with no mutations), and its probability of death—i.e. all of the factors determining its evolutionary success (Dawkins, 1976)—are therefore wholly determined by the local environment in

⁴This mechanism for genome movement was included to compensate for the fact that a genome can only directly affect a single cell at any iteration; at least it can now move from one iteration to the next, and can therefore affect different cells at different times.




Figure 5: EvoCA-B: Evolution of self-stabilising dynamics

which it finds itself. Thus, some local environments will be particularly conducive to the success of a genome, conferring upon it high stability and fecundity, others will be fairly neutral, while others still will be harmful.

In this way, there is a natural selection pressure for genomes that generate local conditions that promote their own survival and reproduction. Genomes can influence their local environments through the action of their genes, but only to a degree—the environment is subject to perturbations from the action of other nearby organisms. Selection pressure for self-generating and self-maintaining dynamics, i.e. autopoiesis (Varela et al., 1974), is therefore an inherent feature of the model.

Example Application: Natural Selection for Genome-Regulated Self-Stabilising Dynamics Typical results for the system as described are shown in Figure 5. These were from a run on a 100x100 grid, seeded with 100 randomly generated genomes placed at random positions. For the first 10000 iterations, fairly random dynamics are observed in Layer 1, as the genomes initiate dynamics in the environment. However, at around iteration 10000, a more dynamically stable pattern begins to appear in a portion of the CA. The pattern, identified by vertical stripes, is not completely static, but is able to regenerate itself under perturbations from the rest of the environment. Over the next 20000 iterations, this pattern spreads to fill the whole space as organisms that possess this phenotype out-compete their rivals. (A movie is available in the supplementary materials.) These dynamics demonstrate the natural selection of genomes that generate dynamics that promote their own survival—in other words, genome-regulated, self-stabilising dynamics.

Discussion

Advantages of the Approach

The results just described illustrate a number of features of the modelling approach being advocated.

Let’s consider the EvoCA-A results first. Success in this

task requires an individual to be sensitive to the signal's presence and location. It also requires that the individual exploits the environmental dynamics for long-range communication, to activate the cells in the goal area. Even in this simple situation, it could be argued that the organisms have evolved sensory and effector apparatus. To say that an individual is responding to a signal rather than just following the "laws of physics" (i.e. the normal CA update rules)—which of course it still is—is justified because the evolutionary selection process has introduced semantics to the signals, from the perspective of the individual, as indications of the task to be performed. During evolution, the successful organisms were selected precisely because they behaved *as if* the state of the cells in the signal position area was a signal, and responded in the appropriate way. Evolution has therefore introduced the potential for a new level of description of the system, where it is more informative to describe the action of an organism in terms of local rules (e.g. organism *A* responds to signal *B* by producing action *C*) rather than in terms of the universal laws of physics (cf. the quote from Pattee in the Introduction section).

Similarly, even though individual genes act by setting the state of single cells in the CA (there is no "glider gene," for example), genomes are able to produce complex actions such as the production of gliders and other patterns of spreading activation. This is because the genes are interacting with the pre-existing dynamics of the environment, by setting initial conditions for those dynamics. Again, these actions acquire semantics from the perspective of the organism through the process of selection during evolution.

The environment in this example is very simple. We could imagine environments with many more possible states per cell, and with much richer dynamics. In such environments, even if an organism's genes could still only directly respond to and activate a limited number of states, we could nevertheless imagine the organism being able to deal with a much wider variety of states, indirectly, by harnessing the environmental dynamics (e.g. by setting up a "chain reaction" to eventually achieve the desired result). In this way, organisms could potentially evolve to exploit almost any property of the environment, even if their genes were still able to perform only a limited subset of actions at the lowest level. Any property or process so incorporated can be expected to be retained if it promotes the evolutionary success of the organism. From this perspective, the evolutionary acquisition of new sensory or effector capabilities is not the problem that it is with other approaches (Dautenhahn et al., 2001).

Another feature of the approach is that phenotypes and abiotic environment are represented as a single system. In cases where we allow multiple organisms to coexist in the environment (e.g. EvoCA-B), organisms are therefore part of the environment experienced by other organisms. This introduces the possibility of rich co-evolutionary dynamics and high evolvability (Waddington, 1969; Odling-Smee

et al., 2003). Evolvability is also increased by the fact that there is no pre-defined specification of the organisation of a phenotype, so this is free to evolve over time.

The approach shifts the focus of the "problem" of evolvability away from the process of self-reproduction (which is taken for granted in appropriate environmental conditions), towards the issue of organism interactions (both organism–organism and organism–environment). It emphasises the view of organisms as self-generating and self-organising organisations, rather than self-reproducing automata.

Potential Criticisms

No modelling approach is perfect, and there are many potential criticisms that could be levelled at EvoCA and at the ideas that it embodies. Some of these are now considered.

Genetic System is Immutable In EvoCA, the perspective of a genome as a source of constraints for a dynamical system is taken to the extreme; genomes play *no* part in the dynamics of the system other than to specify constraints (i.e. they have no material embodiment). This is largely for practical, rather than theoretical, reasons, and means that the design of the system can be kept very simple. The design may be compared to an artificial chemistry; the main difference is this separation of representation of genetic material from the rest of the system. This simplification is not without consequences. It means that an external mechanism is required for interpreting genomes as constraints (this happens at each iteration of the CA), and for writing genomes, with noise, at reproduction (this is performed by the evolutionary algorithm). Another consequence is that the genes are restricted to specifying constraints in a predefined way—in the particular design of EvoCA they are defined to map to the lowest level of the CA dynamics by constraining a specific cell to be in a specific state at a specific time. These restrictions all arise because genomes in EvoCA do not participate in the dynamics of the system at all, except through supplying constraints. This design decision is justified because of the perspective of genomes taken here—that the *fundamental* role of the genome is to supply constraints to the dynamical environment. It should also be noted that it is conceivable that a more complex genotype–phenotype mapping could evolve on top of the given system.

Where are the Organisms? The examples presented, particular for EvoCA-B, are open to the criticism that an organism's phenotype lacks individuality (i.e. there is no recognisable boundary between it and the rest of the environment). The approach takes the view that an organism's phenotype is the set of genome-initiated dynamics in the environment. If it is advantageous for an organism to have a distinct membrane defining its boundary, then this is something we might expect to see evolving. However, particularly if we are modelling the origin of life, we should not assume such a distinction *a priori*. Also note that the question of whether such

a membrane evolves is a question of the *properties of the environment*—is this existence of such a structure possible within the given environment?

The Limitations of Computational Models It could be argued that computational models such as EvoCA are unsuitable for the purpose of studying open-ended evolution, because of their digital nature. Each cell can only exist in one of a small number of discrete states, and therefore the number of states of the system as a whole is limited. A simple counter-argument is that, as the size of the system under consideration grows, the number of possible states soon becomes astronomically large. Furthermore, the complexity of the environment can always be increased, for example by using a multi-layer CA with different dynamics in each layer, or by allowing real-numbered states. More pertinently, we are also interested in dynamics and cycles of states, rather than the state of the system at a single instance—emergent behaviour can of course arise in the dynamics of discrete dynamical systems as well as analogue ones. The crucial point, however, is that the process of evolution, as demonstrated earlier, can endow states with semantic significance from an organism's perspective, at which point it becomes appropriate to describe the system at the level of local rules—the shape of which will depend on the system's specific history—rather than in terms of the underlying laws of physics. More sophisticated arguments have been put forward as to why purely digital devices cannot self-complexify, e.g. (Cariani, 1989). Cariani accepts that “the absence of gradualist pathways [in digital devices] would not preclude evolution entirely” but suggests that “[t]he important point is that purely computational devices do not construct or modify their primitives, and this does foreclose the possibility of fundamental novelty” (*ibid*, p.111). While these arguments hold for the system as a whole (i.e., in the case of EvoCA, the “laws of physics” of the environment), we can still address, given this limitation, how organisms evolve *within* the environment (with the corresponding emergence of semantics already described), starting from very simple forms to progressively control and exploit more and more of the environment's properties. We are not looking for fundamental novelty in the environment itself, but in the way in which organisms interact with it.

Implications

A corollary of this approach is that we can increase the complexity of evolved organisms—while still assuming only a simple set of mechanisms for the evolution of genomes—by increasing the complexity of the abiotic environment. From this perspective, some of the most important research questions are: What features must the environment possess to enable open-ended evolution? What features must it possess to enable the evolution of organisms that we might reasonably regard as *living*? Indeed, how are these two questions related?

(Is one a subset of the other? Are they identical?) We can address these questions not only at the level of features that have long been argued as being necessary for life (e.g. the requirement of a water-like substance in a liquid phase, or the possibility of semi-permeable membranes), but also at the level of fundamental physical properties such as the conservation of matter, energy flow, entropy increase, etc. How critical are each of these features for allowing the possibility of open-ended evolution and the emergence of life? Of course, I am not arguing that this approach should replace traditional ones; rather, it is complementary to them. By taking a different perspective, it highlights the significance of some different questions to consider in our attempt to understand life-as-it-is and life-as-it-could-be.

References

- Bedau, M. A., McCaskill, J. S., Packard, N. H., Rasmussen, S., Adami, C., Green, D. G., Ikegami, T., Kaneko, K., and Ray, T. S. (2000). Open problems in artificial life. *Artificial Life*, 6(4):363–376.
- Cariani, P. (1989). *On the Design of Devices with Emergent Semantic Functions*. PhD thesis, State University of New York at Binghamton.
- Crutchfield, J. P. and Mitchell, M. (1995). The evolution of emergent computation. *Proceedings of the National Academy of Sciences U.S.A.*, 92(23):100742–100746.
- Dautenhahn, K., Polani, D., and Uthmann, T. (2001). Guest editors' introduction: Special issue on sensor evolution. *Artificial Life*, 7(2):95–98.
- Dawkins, R. (1976). *The Selfish Gene*. Oxford University Press.
- Holland, J. H. (1976). Studies of the spontaneous emergence of self-replicating systems using cellular automata and formal grammars. In Lindenmayer, A. and Rozenberg, G., editors, *Automata, Languages, Development*, pages 385–404. North-Holland, New York.
- Maynard Smith, J. (1986). *The Problems of Biology*. Oxford University Press.
- McMullin, B. and Varela, F. J. (1997). Rediscovering computational autopoiesis. In Husbands, P. and Harvey, I., editors, *Fourth European Conference on Artificial Life*, pages 38–47. MIT Press/Bradford Books.
- Odling-Smee, F., Laland, K., and Feldman, M. (2003). *Niche Construction: The Neglected Process in Evolution*. Princeton University Press.
- Pattee, H. (1995a). Artificial life needs a real epistemology. In Morán, F., Moreno, A., Merelo, J., and Chacón, P., editors, *Advances in Artificial Life: Third European Conference on Artificial Life*, LNAI, pages 23–38. Springer.
- Pattee, H. (1995b). Evolving self-reference: Matter, symbols, and semantic closure. *Communication and Cognition—Artificial Intelligence*, 12(1–2):9–28.
- Ray, T. S. (1991). An approach to the synthesis of life. In Langton, C., Taylor, C., Farmer, J., and Rasmussen, S., editors, *Artificial Life II*, volume X of *SFI Studies in the Sciences of Complexity*, pages 371–408. Addison-Wesley.
- Rocha, L. M., editor (2001). *The Physics and Evolution of Symbols and Codes: Reflections on the Work of Howard Pattee*. Special issue of *BioSystems* 60(1–3).
- Taylor, T. (2001). Creativity in evolution: Individuals, interactions and environments. In Bentley, P. J. and Corne, D. W., editors, *Creative Evolutionary Systems*, chapter 1, pages 79–108. Morgan Kaufman.
- Varela, F. J., Maturana, H. R., and Uribe, R. (1974). Autopoiesis: The organization of living systems, its characterization and a model. *BioSystems*, 5:187–196.
- von Neumann, J. (1966). *The Theory of Self-Reproducing Automata*. University of Illinois Press, Urbana, Ill.
- Waddington, C. (1969). Paradigm for an evolutionary process. In Waddington, C., editor, *Towards a Theoretical Biology*, volume 2, pages 106–128. Edinburgh University Press.
- Yaeger, L. (1994). Computational genetics, physiology, metabolism, neural systems, learning, vision and behavior or poly-world: Life in a new context. In Langton, C., editor, *Proceedings of Artificial Life III*, pages 263–298. Addison-Wesley.