

On Self-Reproduction and Evolvability^{*}

Tim Taylor

Institute of Perception, Action and Behaviour,
Division of Informatics, University of Edinburgh,
5 Forrest Hill, Edinburgh EH1 2QL, Scotland.
tim.taylor@ed.ac.uk
<http://www.dai.ed.ac.uk/daidb/homes/timt/>

Abstract. Von Neumann's architecture for self-reproducing, evolvable machines is described. From this starting point, a number of issues relating to self-reproduction and evolution are discussed. A summary is given of various arguments which have been put forward regarding the superiority of genetic reproduction over self-inspection methods. It is argued that programs in artificial life platforms such as Tierra reproduce genetically rather than by self-inspection (as has previously been claimed). However, the distinction is blurred because significant parts of the reproduction process in Tierran programs are implicitly encoded in the Tierran operating system. The desirable features of a structure suitable for acting as a seed for an open-ended evolutionary process are discussed. It is found that the properties of such a structure are somewhat different to those of programs in Tierra-like platforms. These analyses suggest ways in which the evolvability of individuals in artificial life platforms may be improved, and also point to a number of open questions.

1 Introduction

In the late 1940s and early 1950s, John von Neumann devoted considerable time to the question of how complicated machines could evolve from simple machines.¹ Specifically, he wished to develop a formal description of a system that could support self-reproducing machines which were robust in the sense that they could withstand some types of mutation and pass these mutations on to their offspring. Such machines could therefore participate in a process of evolution.

Inspired by Alan Turing's earlier work on universal computing machines [3], von Neumann devised an architecture which could fulfil these requirements. The machine he envisaged was composed of three subcomponents [2]:

1. A general *constructive* machine, **A**, which could read a description $\phi(\mathbf{X})$ of another machine, **X**, and build an instance of **X** from this description:

$$\mathbf{A} + \phi(\mathbf{X}) \rightsquigarrow \mathbf{X} \tag{1}$$

^{*} This paper is an abbreviated version of certain sections of [1].

¹ Von Neumann had difficulties in defining precisely what the term 'complicated' meant. He said "I am not thinking about how involved the object is, but how involved its purposive operations are. In this sense, an object is of the highest degree of complexity if it can do very difficult and involved things." [2].

(where + indicates a single machine composed of the components to the left and right suitably arranged, and \rightsquigarrow indicates a process of construction.)

2. A general *copying* machine, **B**, which could copy the instruction tape:

$$\mathbf{B} + \phi(\mathbf{X}) \rightsquigarrow \phi(\mathbf{X}) \quad (2)$$

3. A *control* machine, **C**, which, when combined with **A** and **B**, would first activate **B**, then **A**, then link **X** to $\phi(\mathbf{X})$ and cut them loose from $(\mathbf{A} + \mathbf{B} + \mathbf{C})$:

$$\mathbf{A} + \mathbf{B} + \mathbf{C} + \phi(\mathbf{X}) \rightsquigarrow \mathbf{X} + \phi(\mathbf{X}) \quad (3)$$

Now, if we choose **X** to be $(\mathbf{A} + \mathbf{B} + \mathbf{C})$, then the end result is:

$$\mathbf{A} + \mathbf{B} + \mathbf{C} + \phi(\mathbf{A} + \mathbf{B} + \mathbf{C}) \rightsquigarrow \mathbf{A} + \mathbf{B} + \mathbf{C} + \phi(\mathbf{A} + \mathbf{B} + \mathbf{C}) \quad (4)$$

This complete machine plus tape, $[\mathbf{A} + \mathbf{B} + \mathbf{C} + \phi(\mathbf{A} + \mathbf{B} + \mathbf{C})]$, is therefore self-reproducing. From the point of view of the evolvability of this architecture, the crucial feature is that we can add the description of an arbitrary additional automaton **D** to the input tape. This gives us:

$$\mathbf{A} + \mathbf{B} + \mathbf{C} + \phi(\mathbf{A} + \mathbf{B} + \mathbf{C} + \mathbf{D}) \rightsquigarrow \mathbf{A} + \mathbf{B} + \mathbf{C} + \mathbf{D} + \phi(\mathbf{A} + \mathbf{B} + \mathbf{C} + \mathbf{D}) \quad (5)$$

Furthermore, notice that if the input tape $\phi(\mathbf{A} + \mathbf{B} + \mathbf{C} + \mathbf{D})$ is mutated in such a way that the description of automaton **D** is changed, but that of **A**, **B** and **C** are unaffected (that is, the mutated tape is $\phi(\mathbf{A} + \mathbf{B} + \mathbf{C} + \mathbf{D}')$), then the result of the construction will be:

$$\begin{aligned} \mathbf{A} + \mathbf{B} + \mathbf{C} + \mathbf{D} + \phi(\mathbf{A} + \mathbf{B} + \mathbf{C} + \mathbf{D}) &\xrightarrow{\text{mutation}} \\ \mathbf{A} + \mathbf{B} + \mathbf{C} + \mathbf{D}' + \phi(\mathbf{A} + \mathbf{B} + \mathbf{C} + \mathbf{D}') &\end{aligned} \quad (6)$$

The reproductive capability of the architecture is therefore robust to some mutations (specifically, those mutations which only affect the description of **D**), so the machines are able to evolve. Von Neumann pointed out that the action of the general copying automaton, **B**, was the decisive step which gave his architecture the capacity for evolving machines of increased complexity, because **B** is able to copy the description of any machine, no matter how complicated [2] (p.121). This ability is clearly demonstrated in Reaction 5 above.

2 General Issues of Reproduction

The major focus of this paper is self-reproduction in the specific context of evolution. However, before continuing it is useful to briefly consider some more general issues relating to reproduction.

When looking at any sort of reproduction, it is helpful to look at the process by which reproduction is accomplished from a number of different perspectives. Two important ones are:

1. The degree to which the reproduction process is *explicitly encoded* on the configuration being reproduced, rather than being *implicit in the physical laws* of the world.
2. The number of different configurations that exist, connected by mutational pathways, that are capable of reproducing their specific form (i.e. the distinction between *limited hereditary* reproducers and *indefinite hereditary* reproducers). From the point of view of an individual reproducer, this can be expressed in terms of the proportion of all possible mutations it may experience that will result in the production of distinct, yet viable, reproducers.

There are a number of points to note about these distinctions. First it should be said that (2), in contrast to (1), does not properly relate to individual reproducers *per se*, but rather to lineages of reproducers. It is therefore not relevant when considering self-reproduction in and of itself, but is an important factor when considering the evolutionary potential of a class of reproducers.

Secondly, the two distinctions are generally independent of each other, although the more explicitly encoded the reproduction algorithm is, the less likely, in general, it is to be an indefinite hereditary reproducer (because of the increased chance of mutations disrupting the copying process; see Section 3.3).

3 Self-Reproduction and Open-Ended Evolution

I now wish to return to issues of reproduction in the specific context of evolution. In this section I will concentrate on a number of these issues in turn.

3.1 Trivial versus Non-Trivial Reproduction

Notice that in much of the recent artificial life work with self-reproduction (e.g. [4]), the distinction between trivial and non-trivial self-reproduction is perceived to be a distinction on the implicit-explicit axis.² However, from an evolutionary point of view, the limited-indefinite heredity axis is clearly the most relevant. Indeed, this is exactly what von Neumann himself says: “One of the difficulties in defining what one means by self-reproduction is that certain organizations, such as growing crystals, are self-reproductive by any naive definition of self-reproduction, yet nobody is willing to award them the distinction of being self-reproductive. A way around this difficulty is to say that self-reproduction includes the ability to undergo inheritable mutations as well as the ability to make another organism like the original” [2].

Barry McMullin has presented an enlightening discussion on the history of the confusion over von Neumann’s work, which he refers to as the ‘von Neumann Myth’ (see, for example, Section 4.2.7 in [5]). One result of this confusion has

² From this point of view, an example of trivial reproduction in a cellular automata space would be where the state of a single cell is reproduced in neighbouring cells purely due to the CA’s transition rules.

been that the majority of subsequent research concerning this issue of trivial self-reproduction has concentrated on the implicit-explicit distinction, rather than the limited-indefinite heredity distinction.

Von Neumann's work on self-reproduction concerned the question of how machines might be able to *evolve increased complication* in order to perform increasingly complex tasks. This is why his design for a self-reproducing machine had to be capable of universal construction, and why it was designed in such a way that it could withstand some kinds of mutations.

3.2 Genetic Reproduction versus Self-Inspection

Von Neumann's architecture was designed specifically to allow for a possible increase in complexity and efficiency of machines by evolution. However, even if we accept that his design *is* a solution to this problem, it is by no means the only conceivable solution, as von Neumann himself was well aware. In particular, he also discussed the possibility of a machine which built a copy of itself by actively inspecting its parts, without the need for this design information to be duplicated on a tape (i.e. without a 'genetic' description). Indeed, systems which reproduce by self-inspection have been designed by Laing, e.g. [6], and by Ibáñez and colleagues [7]. From the point of view of designing artificial life systems, we would like to know which of the possible architectures we should employ (according to factors such as their relative simplicity, efficiency, etc.).

Although he certainly did not *prove* that reproduction by self-inspection could not support open-ended evolution, von Neumann did suggest a number of reasons why his genetic architecture would be a more powerful and more general design for this purpose. First of all, as mentioned in Section 1, he noted that the essential feature which allowed his automata to overcome the otherwise seemingly valid rule that machines are necessarily superior (in size and in organization) to their output, was that they contained a general copying automaton \mathbf{B} , which was capable of copying any linear tape [2] (p.121). Although \mathbf{B} is of fixed, finite size, it is able to copy a tape of any size. Now, this action of copying a tape is essentially reproduction by self-inspection, but this is generally a straightforward task for a linear tape. The major problems arise when trying to copy a two- or three-dimensional structure by the same method, for example in specifying the precise spatial relationships between parts, and in unfolding multidimensional forms. Von Neumann also pointed out that self-inspection requires that we have a representation which is 'quasi-quiescent' in the sense that it can be read (for the purposes of copying and possibly for interpretation) without being essentially disturbed. With a separate genetic description, we only require that this description is quasi-quiescent, but copying by self-inspection would require that the whole structure to be copied would have this quasi-quiescent property. In general, however, most machines would not have this property, nor would we want to restrict ourselves to only considering those machines which did. In conclusion, von Neumann says: "To sum up, the reason to operate with 'descriptions' . . . instead of the 'originals' . . . is that the former are quasi-quiescent (i.e. unchanging, not in an absolute sense, but for the purposes of the exploration

that has to be undertaken), while the latter are live and reactive. In the situation in which we are finding ourselves here, the importance of descriptions is that they replace the varying and reactive originals by quiescent and (temporarily) unchanging semantic equivalents and thus permit copying. Copying, as we have seen above, is the decisive step which renders self-reproduction (or, more generally, reproduction without degeneration in size or level of organization) possible” [2] (p.122–123).

From a biological perspective, Waddington has made the same point. While discussing possible reasons for the universal adoption of genetic architectures for self-reproduction by biological life, he suggested that the issue “is presumably related to the problem [of] how to combine a store which is unreactive enough to be reliable, with something which interacts with the environment sufficiently actively to be ‘interesting’” [8] (p.118).

McMullin has pointed out that von Neumann’s genetic architecture also effectively decouples the geometry of the variational space of the reproducers (i.e. the space of the genetic tapes) from the peculiarities of the environment in which they exist (i.e. the space of the phenotype) [5] (pp.191–193). In addition, recall from Section 1 that the architecture will accept any tape of the general form $\phi(\mathbf{A} + \mathbf{B} + \mathbf{C} + \mathbf{D})$. Assuming the the description of \mathbf{D} on the tape can be separated from the description of \mathbf{A} , \mathbf{B} and \mathbf{C} ,³ this design guarantees that mutations which affect the part of the tape describing automaton \mathbf{D} will not interfere with the reproductive capacity of the machine. Machines which reproduce by self-inspection would generally not have this localisation property. This being the case, we would not always be able to say that there was a particular section of the machine which could be disrupted by mutation without interfering with the machine’s ability to reproduce.

It is interesting to ask whether programs in artificial life platforms such as Tierra [10] reproduce according to von Neumann’s genetic architecture or rather by self-inspection. As the arguments of the previous paragraphs suggest that a marked difference exists in the evolutionary potential of these two methods, it is an important question, but it has not received much discussion in the literature. McMullin argued that these programs are reproducing by self-inspection [5] (p.200). Ibáñez and colleagues appear to agree [7] (p.574). In contrast, I would like to suggest that they can sensibly be analysed in terms of von Neumann’s genetic architecture. I do not have adequate room to argue the case fully here,

³ This is not an inherent property of the architecture *per se*, but von Neumann’s analysis of evolvability did assume a ‘compositional’ structure in the language of the tape descriptions (see Section 1). His cellular automata model [2], and Pesavento’s recent implementation of a very similar design [9], are existence proofs that it is possible to build a self-reproducing automaton with such a compositional genetic structure. Interestingly, however, despite his design for the cellular automata model, von Neumann also argued that “it is better not to use a description of the pieces and how they fit together, but rather a description of the consecutive steps to be used in building the automaton” [2]. In other words, the information should be in the form of a developmental ‘recipe’ rather than a ‘blueprint’. Further discussion of this topic can be found in Chapter 7 of [1].

so the following paragraphs are included merely for the purpose of stimulating discussion on the subject.

Before I begin, I would like to make a couple of general points, which might help to reorient the reader to my perspective. Firstly, I believe that the notion of a phenotype fundamentally involves *interaction* with the environment (and that this is the essential distinction between the notions of phenotype and genotype—the latter being an informational concept). When I talk about phenotypes in the following, therefore, and specifically when I talk about the automata **A**, **B**, **C** and **D**, I am interested in the role these phenotypic structures play—their function—rather than the details of implementation or of how that function is achieved. Secondly, note that the terminology commonly used to describe reproducers in Tierra-like systems is somewhat different to that used for von Neumann’s work. Because of the similarity between Tierra-like operating systems and those of standard digital computers, the actions of Tierran reproducers are often referred to as computations rather than constructions, even when a reproducer is in the process of building a new copy of itself. However, this process of reproduction is, of course, central to the Tierra approach, and I believe that this procedure of building a copy of a program in a different part of memory is, in all the relevant details, a process of construction in just the same way as construction processes in von Neumann’s cellular automata model. In the following, also remember that von Neumann’s general constructing automaton **A** is the machinery which *interprets* the tape to produce a new machine (phenotype), and the general copying automaton **B** copies the tape uninterpreted.

At first sight it might seem that there is no separate genetic description of the program in a Tierra-like system. The picture is complicated by the fact that the machinery which interprets the program (i.e. automaton **A**) does not reside in the same part of the computer in which the program itself is stored. The state information for this machinery—a program’s ‘virtual CPU’ (i.e. the instruction pointer, stacks, registers, etc.)—is generally represented in an independent area of memory to the program’s instructions. Furthermore, the actual ‘interpreting machinery’ of the virtual CPU is encoded in the global operating system provided by the platform, and is in this sense implicit in the program’s environment. Additionally, the control automaton **C**, which controls when the instructions in the program get executed, is also implicit in the part of the operating system which governs mechanisms such as how a program’s instruction pointer is updated after the execution of each instruction. All that is left to be explicitly encoded by the program, therefore, is the copying automaton **B**, and potentially any other arbitrary automaton **D**.

Now, the instructions which make up the program exist in an unreactive state in the system’s random-access memory. It is only when the control automaton **C** transfers instructions to the interpreting automaton **A** that they become ‘active’. Looked at in this way, we can see that it is the *behaviour* of the program (including looping, jumping, etc.) that is the result of automaton **A** interpreting the unreactive genetic description. This behaviour is therefore the equivalent to the constructed machine (and the actions it performs—i.e., the phenotype) in

von Neumann's design, and the string of instructions residing in the random-access memory (which is normally referred to as the program) is the tape or genetic description of this phenotype. It is perhaps easier to see the distinction if one considers a parallel program, with multiple processes (with different state information) using the same program listing.

I therefore suggest that a self-reproducing program in a Tierra-like system is consistent with von Neumann's architecture. However, as automata **A** and **C** are largely implicit in the environment in which the programs reside (the only explicit representation being the state information in a program's virtual CPU), and are certainly not encoded by the individual programs, we can see that a 'program', in the sense of a string of instructions in the system's random-access memory, corresponds to the tape $\phi(\mathbf{B} + \mathbf{D})$ in von Neumann's scheme.

The situation is complicated not only because the interpretation machinery resides partly implicitly in the environment, and partly in a different area of memory, but also for (at least) one further reason. I am claiming that the string of instructions comprising the 'program' in random-access memory should be viewed as the genetic tape in a von Neumann style self-reproduction architecture. Now, von Neumann pointed out that the process of copying the tape in his automaton was essentially itself a process of self-inspection. In this sense, Tierran programs do reproduce by self-inspection. However, the overall mechanism for reproduction, including the implicit encodings of the interpretation and control automata, fits in with von Neumann's architecture, in which the copying of the tape by self-inspection is an integral feature. The major consequence of this is that programs in Tierra-like systems should, all else being equal, have similar evolutionary potential to von Neumann's self-reproducing automata, because extra instructions can be added to the end of the 'tape' and subject to mutations. As long as the mutations do not affect that part of the tape which encodes the self-reproduction algorithm, they will be inherited without disrupting the capacity of the program to reproduce.

3.3 Implicit versus Explicit Encoding

The preceding arguments have led us to consider the question of implicit versus explicit encoding of automata. However, rather than the general question that has been the subject of much debate relating to trivial versus non-trivial reproduction, here we are interested in rather more specific questions relating to von Neumann's architecture. Now, as we are interested in the evolution of these self-reproducing machines, and as the inheritable information of each machine (i.e. the part which gets passed on from parent to offspring) is contained on the tape ϕ , I will assume that the tape must be explicitly represented in some fashion, otherwise there would be nothing which could evolve. We can now ask *which parts* of the $[\mathbf{A} + \mathbf{B} + \mathbf{C} + \mathbf{D}]$ architecture are explicitly encoded on the tape ϕ , and which are implicit in the environment. Of course, even the behaviour of those parts which are represented on the tape will still to some extent be encoded in the 'laws of physics' of the environment, but I think the analysis is nevertheless worthwhile.

Considering von Neumann’s architecture for a self-reproducing automaton, it is clear that all four subcomponents, **A**, **B**, **C** and **D**, are very explicitly encoded on the tape $\phi(\mathbf{A} + \mathbf{B} + \mathbf{C} + \mathbf{D})$; the environment in which the automaton exists implicitly encodes only very low-level actions in the form of the local transition rules of individual cells. The analysis of self-reproducing programs in Tierra-like systems above suggests that in these systems, **B** and **D** are explicitly encoded on the tape $\phi(\mathbf{B} + \mathbf{D})$, but **A** and **C** are implicitly encoded in the environment (the operating system). Notice that with this design the ‘genetic code’ which maps the genotype $\phi(\mathbf{B} + \mathbf{D})$ to the phenotype $[\mathbf{B} + \mathbf{D}]$ cannot itself evolve, because the interpretation automaton **A** is not encoded on the tape.

It is interesting to speculate on what information we might desire to be explicitly encoded on a structure which would be suitable for acting as a robust initial seed for an open-ended evolutionary process. I will refer to such a structure as ‘proto-DNA’. Now, we would like our proto-DNA to be an indefinite hereditary replicator if it is to be such a seed. In other words, it should be able to exist in an unlimited number of configurations which retain the ability to reproduce. If the copying process is encoded on the tape itself, then mutations have the potential to disrupt its ability to be reproduced. It would therefore seem desirable that the copying automaton **B** of our proto-DNA be largely implicitly encoded in the environment. Note that this would not necessarily prevent a more complicated, and possibly more reliable, explicit copying process **B’**—genetically encoded as $\phi(\mathbf{B}')$ —later evolving from (but still based upon) the simpler implicit process, as indeed seems to have happened during biological evolution.

If the copying procedure for our proto-DNA is implicitly encoded in the environment, however, any configuration of proto-DNA would, all else being equal, be able to reproduce as well as any other. In other words, there would be no basis for preferentially selecting some configurations over others, and therefore no basis for an evolutionary process. Specific configurations of proto-DNA must therefore have some specific properties that are selectively significant. Models of the origin of life commonly presume that these simple phenotypic properties were things such as increased stability of the molecule, simple control of the local environment, catalytic activity, etc. (e.g. [11], [12], [13]).

At the initial stages of an evolutionary process, however, we would not expect there to be mechanisms for explicitly decoding the proto-DNA; in other words, the interpretation machinery **A** is implicit. This means that particular configurations of proto-DNA should have some specific phenotypic properties (such as the ability to act as catalysts) which can be determined directly from their structure rather than having to be explicitly decoded from the genotype. We could therefore regard the proto-DNA as merely $\phi(\mathbf{D})$, meaning that particular configurations have particular phenotypes associated with them, which are (a) not related to the process of self-reproduction *per se*, and (b) do not require to be decoded by an explicit interpretation automaton **A**.⁴ Regarding the kinds of

⁴ I am also assuming here that the domain of interaction of these phenotypes is within the environment shared by the evolving population (i.e. the phenotypes can act upon other biotic or abiotic components in the environment). This is in contrast

simple phenotypes that we might wish to be available to our proto-DNA, the possibilities seem endless. Graham Cairns-Smith observes: “It is almost too easy to imagine possible uses for phenotype structures—because the specification for an effective phenotype is so sloppy. A phenotype has to make life easier or less dangerous for the genes that (in part) brought it into existence. There are no rules laid down as to how this should be done” [12] (p.106). If more complicated phenotypes are to arise later on in the evolutionary process, however, we require that the proto-DNA at least has the potential for explicit interpretation machinery \mathbf{A}' and control machinery \mathbf{C}' to become associated with it. This would involve some form of specific reaction to subsections of information in the proto-DNA, but more work is needed to fully identify how this potential for explicit interpretation might be assured.

4 Discussion

It has been argued that programs in artificial life platforms such as Tierra conform to von Neumann’s genetic architecture for self-reproducing machines. Specifically, the listing of such a program corresponds to the tape $\phi(\mathbf{B} + \mathbf{D})$. In contrast, the analysis of the desirable properties of proto-DNA (i.e. a class of object capable of acting as a seed for an open-ended evolutionary process) suggests that such an object would correspond to the tape $\phi(\mathbf{D})$ only. The fact that the copying process \mathbf{B} is explicitly encoded in Tierran programs means that it is susceptible to disruption by mutations and perturbations from the environment. This is why the interactions between programs in platforms such as this have to be restricted—in Tierra, for example, direct interaction between programs is restricted to the reading by one program of the instructions of its neighbours.

It is likely that many of the more interesting ecological and evolutionary phenomena in the biosphere arise because organisms are able to interact in much richer ways. Most importantly, biological organisms are embedded in a material world, and therefore represent useful resources of matter and energy for potential use by other organisms. Without such an unrestricted range of allowable interactions, and without the system being grounded on a material basis (i.e. where organisms are composed of structural units which are, at their lowest level, conserved, and which are in limited supply), it is doubtful whether any selection pressure can exist for organisms to evolve properties such as self-maintenance. Also, it is only with such a material grounding that ecological phenomena such as food webs and trophic levels can be realised. If we wish to allow artificial life models the capacity to evolve in these ways, we must model a material environment, and allow the individual organisms much more freedom in their interactions. However, if we were to model organisms as self-reproduction algorithms (as in Tierra) in such an environment, they would prove very brittle, because the explicitly-encoded copying process \mathbf{B} would be very easily disrupted.

to models such as genetic algorithms, where the replicators do not directly interact with other replicators and selection is determined by an extrinsic fitness function (thereby limiting the potential for open-ended evolution).

The advantage of the proposed proto-DNA structure, where **B** is implicit in the environment, is that we can relax the restrictions on interactions between self-reproducing machines while maintaining the robustness of the individuals. Such a system would therefore have the potential for evolving much richer ecologies and symbioses. This approach brings with it many new issues which will have to be addressed, such as what range of phenotypes should be available to the proto-DNA structures, and how to ensure that they have the potential for evolving more explicit interpretation machinery to encode more complicated phenotypes. The approach therefore involves a shift of focus away from the process of self-reproduction *per se*, towards questions relating to phenotypes, the environment, and interactions between organisms. Progress in these directions might give us a better insight into the essential requirements for biological evolution, and might also allow us to build artificial life models with improved evolutionary potential.

Acknowledgements

Thanks to my PhD supervisor, John Hallam, and to the anonymous reviewers, for their useful comments and criticisms. Thanks also to EPSRC for providing financial support.

References

1. Taylor, T.: From Artificial Evolution to Artificial Life. PhD thesis, Division of Informatics, University of Edinburgh (1999)
2. von Neumann, J.: The Theory of Self-Reproducing Automata. University of Illinois Press, Urbana, Ill. (1966)
3. Turing, A.M.: On computable numbers, with an application to the entscheidungsproblem. Proc. London Mathematical Soc., Series 2 **42** (1936) 230–265
4. Langton, C.G.: Self-reproduction in cellular automata. Physica D (1984) 135–144
5. McMullin, B.: Artificial Knowledge: An Evolutionary Approach. PhD thesis, Department of Computer Science, University College Dublin (1992)
6. Laing, R.: Automaton models of reproduction by self-inspection. Journal of Theoretical Biology **66** (1977) 437–456
7. Ibáñez, J., Anabitarte, D., Azpeitia, I., Barrera, O., Barrutieta, A., Blanco, H., Echarte, F.: Self-inspection based reproduction in cellular automata. In Morán, F., Moreno, A., Merelo, J., Chacón, P., eds.: Third European Conference on Artificial Life, Springer (1995) 564–576
8. Waddington, C.: Paradigm for an evolutionary process. In Waddington, C., ed.: Towards a Theoretical Biology. Volume 2. Edinburgh Univ. Press (1969) 106–128
9. Pesavento, U.: An implementation of von Neumann's self-reproducing machine. Artificial Life **2** (1995) 337–354
10. Ray, T.S.: An approach to the synthesis of life. In Langton, C., Taylor, C., Farmer, J., Rasmussen, S., eds.: Artificial Life II. Addison-Wesley, Redwood City, CA (1991) 371–408
11. Eigen, M., Schuster, P.: The hypercycle: A principle of natural self-organization. Die Naturwissenschaften **64** (1977) 541–565
12. Cairns-Smith, A.: Seven Clues to the Origin of Life. Cambridge Univ. Press (1985)
13. Szathmáry, E., Demeter, L.: Group selection of early replicators and the origin of life. Journal of Theoretical Biology **128** (1987) 463–486