John von Neumann and the Evolutionary Growth of Complexity: Looking Backwards, Looking Forwards...

Barry McMullin
http://www.eeng.dcu.ie/~mcmullin/

©2000

Presented at
Artificial Life VII
August 1–6, 2000, Portland, Oregon.

Dublin City University
Research Institute for Networks and Communications Engineering
Artificial Life Laboratory
Technical Report Number: bmcm-2000-01

Abstract

In the late 1940's John von Neumann began to work on what he intended as a comprehensive "theory of [complex] automata". He started to develop a book length manuscript on the subject in 1952. However, he put this aside in 1953, apparently due to pressure of other work. Due to his tragically early death in 1957, he was never to return to it. The draft manuscript was eventually edited, and combined for publication with some related lecture transcripts, by Burks (1966). It is clear from the time and effort which von Neumann invested in it that he considered this to be a very significant and substantial piece of work. However: subsequent commentators (beginning even with Burks) have found it surprisingly difficult to articulate this substance. Indeed, it has since been suggested that von Neumann's results in this area are either trivial, or, at the very least, could have been achieved by much simpler means. It is an enigma. In this paper I review the history of this debate (briefly) and then present my own attempt at resolving the issue by focusing on an analysis of von Neumann's *problem situation* (Popper, 1976). I claim that this reveals the true depth of von Neumann's achievement and influence on the subsequent deveopment of this field; and, further, that it generates a whole family of new consequent problems which can still serve to inform—if not actually define—the field of Artificial Life for many years to come.

1 Burks' Problem: Machine Self-Reproduction

> ... This result is obviously substantial, but to express its real force we must formulate it in such a way that it cannot be trivialized...
>
> —Burks (1970*a*, p. 49)

Arthur Burks makes this comment following a 46 page recapitulation of John von Neumann's design for a self-reproducing machine (realised as a configuration embedded in a 29-state, two dimensional, cellular automaton or CA). The comment can be fairly said to have initiated an extended debate about the significance of this work of von Neumann's, which has waxed and waned over time, but still persists to the present day.

Von Neumann's design is large and complex, and relies for its operation on exact and intricate interactions between the many relatively simple parts. In that sense, it is certainly substantial; but Burks is absolutely accurate in pointing out that this intricacy, in itself, is not necessarily interesting or significant. In particular, if the same "results" could be achieved, or the same "problem" solved, with drastically simpler machinery, then the interest of von Neumann's design would be critically undermined.

This is no idle concern on Burks' part. As he himself points out, within the CA framework, one can easily formulate a simple rule whereby a cell in a distinct state (labelled, say, `1`) will cause adjacent quiescent cells (labelled, say, `0`) to transit to state `1` also. By essentially the same definition or criterion as was applied to von Neumann's system, such a single cell, state `1`, configuration, would qualify as a self-reproducing machine—and would seem to render von Neumann's fabulously baroque design completely redundant.

Burks concluded that "... what is needed is a requirement that the self-reproducing automaton have some minimal complexity." And at first sight, this does seem eminently reasonable. Presumably (?) it is relatively easy to construct a "simple" machine, by whatever means; therefore it need not surprise us unduly if we manage to concoct a "simple" machine which can construct other "simple" machines, including ones "like" itself, and which therefore qualifies as self-reproducing. Whereas, it is relatively difficult to construct a "complex" machine, by any means; and therefore it may well be a challenging problem to exhibit a "complex" machine that is capable of self-reproduction. Von Neumann's machine certainly appears "complex", and certainly succeeds in constructing other machines like itself (i.e., in reproducing itself). So if we could just more formally express the precise sense in which von Neumann's machine *is* "complex", then we might indeed be able to clarify the "real force" of his achievement.

However, even at this point, we should be at least somewhat wary—because, while von Neumann himself certainly did introduce and discuss the notion of "complexity" in relation to this work, he did *not* attempt any formalisation of it. Indeed, he described his own concept of complexity as "vague, unscientific and imperfect" (von Neumann, 1949, p. 78). It would therefore seem unlikely that the significance of his eventual results should actually *rely* on such a formalisation.

Nonetheless, Burks went on to propose just such a formal criterion of complexity—namely the ability to carry out universal computation. And by this criterion, von Neumann's design (or, at least, a straightforward derivative of it) would qualify as a "complex" self-reproducing machine, and thus be clearly distinguished from the single cell, "`1` state self-reproducer," which would remain merely a "simple" (and thus *trivial*) self-reproducing machine.

This seems a reasonable enough suggestion by Burks; though I would emphasise again that, as far as I have been able to establish, such a thing was never proposed by von Neumann himself—and, indeed, it jars seriously with von Neumann's calculated refusal to formalise complexity.

In any case, it turns out that Burks' suggestion is unsatisfactory and unsustainable. While it is true that von Neumann's machine (suitably formulated) can satisfy Burks' criterion for "complex" self-reproduction, this still represents an interesting result only if this criterion cannot be satisfied by very much simpler means. But in fact—and with hindsight, this now seems unsurprising—Burks' criterion *can* be satisfied by much simpler means than those deployed by von Neumann. This is because universal computation, *per se,* does not actually require particularly complex machinery (see, for example, Minsky, 1967).

This fact was formally demonstrated by Herman (1973), when he essentially showed how the basic single cell, `1` state, self-reproducer described earlier, could be combined with a single cell capable of universal computation. This results in a CA system in which the individual cells are "simpler" than in von Neumann's CA (i.e., have fewer states), and yet there are single cell configurations capable of both self-reproduction and universal computation. Granted, the universal computation ability relies on an adjacent, indefinitely long, "tape" configuration;

but that was equally true of the universal computation ability of von Neumann's design, and is not a relevant distinguishing factor.

Herman draws the following conclusion (Herman, 1973, p. 62):

> ... the existence of a self-reproducing universal computer-constructor in itself is not relevant to the problem of biological and machine self-reproduction. Hence, there is a need for new mathematical conditions to insure non-trivial self-reproduction.

So we see that Herman rejects Burks' specific criterion, while still continuing to accept Burks' formulation of the *issue* at stake—namely the identification of a suitable criterion for distinguishing "non-trivial" self-reproduction, albeit in Herman's version this is no longer explicitly tied to the notion of "complexity".

The discussion was taken up again by Langton (1984) (though apparently without reference to Herman's work). He presented a rather different analysis of Burks' criterion, but with ultimately complementary results. Langton pointed out that, as a general principle, there is little evidence to suggest that living organisms contain universal computation devices embedded within them. Since the self-reproduction of living organisms is presumably to be regarded as non-trivial (by definition–at least in this context), we should not, therefore, adopt universal computation as a criterion. So in this respect, albeit for different reasons, Langton concurs with Herman.

More importantly, Langton goes on to suggest a specific alternative criterion. He points out that self-reproduction in living organisms relies on a decomposition of the organism into two parts or components playing very distinct roles in the reproductive process:

1. The *genotype*, being an informational pattern stored or recorded in some sort of quasi-static or stable carrier. This information is transcribed or *copied* into a corresponding carrier in the offspring.

2. The *phenotype*, being the visible, active, dynamic, interacting part of the organism. The phenotype of the offspring is created by some sort of *decoding* or *interpretation* of the genotype (rather than by a copying of the parental phenotype).

Langton does not explain just *why* such a decomposition may be important, but rather seems to accept its pervasiveness among biological organisms as reason enough to adopt it as a criterion. And, in some sense it "works", because, indeed, von Neumann's self-reproducing design does have this architecture, whereas (say) Herman's self-declared "trivial" design does not. So it seems like this may be a satisfactory or even illuminating demarcation.

However: Langton did not stop there. With this new criterion in hand, he went on to consider whether it could be satisfied with a design significantly simpler than von Neumann's—and it transpires that it *can*. In fact, Langton was able to present a design for a CA space which itself is rather simpler than von Neumann's (i.e., having fewer states per cell), into which he could embed a self-reproducing automaton which, like von Neumann's, has an explicit decomposition into genotype and phenotype, but which is vastly smaller and simpler, occupying a region of just 150 cells—compared to the several hundred thousand cells required for von Neumann's device!

Langton's automaton is certainly still quite intricate, and its design involved a subtle interplay between designing the CA itself and designing the automaton to be embedded within it. In this sense it is a significant and substantive achievement. But is remains very unsatisfactory from the point of view of evaluating von Neumann's work. If Langton's criterion for non-trivial self-reproduction is accepted as an appropriate measure to judge von Neumann's work by, then we must conclude that the latter's design is vastly more complex than necessary. While this *may* be true, I suggest that we should be reluctant to accept it without some much more substantive rationale for Langton's criterion. Or to put it another way, there may still be more "real force" to von Neumann's achievement than is captured or implied by Langton's criterion.

2 Von Neumann's Problem: The Evolutionary Growth of Complexity

I propose to resolve this enigma in a rather different way.

Firstly, I fully agree with Burks that to appreciate the full force of von Neumann's work, we must understand what *problem* he was attempting to solve. In particular, if it should turn out that this problem can be solved by trivial (Herman) or at least much simpler (Langton) means then we should have to conclude that it was not such a substantial achievement after all. Where I differ from these, and indeed, most other, commentators, is that I think it is

a mistake to view von Neumann's problem as having been wholly, or even largely, concerned with *self-reproduction*!

Of course, this is not to deny that von Neumann did, indeed, present a design for a self-reproducing automaton. I do not dispute that at all. Rather, my claim is that this self-reproducing capability, far from being the object of the design, is actually an incidental—indeed, *trivial*, though highly serendipitous—corollary of von Neumann's having solved at least some aspects of a far deeper problem.

This deeper problem is what I call the *evolutionary growth of complexity*. More specifically, the problem of how, in a general and open-ended way, machines can manage to construct other machines more "complex" that themselves. For if our best theories of biological evolution are correct, and assuming that biological organisms are, in some sense, "machines", then we must hold that such a constructive increase in complexity has happened not just once, but innumerable times in the course of phylogenetic evolution. Note that this claim does not rely on any sophisticated, much less formal, definition of complexity; it requires merely the crudest of qualitative rankings. Nor does it imply any *necessary* or consistent growth in complexity through evolution, but merely an acceptance that complexity has grown dramatically in *some* lineages.

Why is this growth of complexity a problem? Well, put simply, all our pragmatic experience of machines and engineering points in the opposite direction. In general, if we want to construct a machine of any given degree of complexity, we use even more complex machinery in its construction. While this is not definitive, it is surely suggestive of a difficulty.

To make all this a little more concrete, imagine that we could exhibit the following:

- Let there be a large (ideally, infinite) class of machine "types" or "designs"; call this M.
- Now consider those machine types, which we shall call *constructors*, which are capable of constructing *some* other (types of) machine. For any given $m \in M$ let $O(m)$ ("offspring of m") denote that subset of M which m is capable of constructing. m is then a constructor precisely provided $O(m) \neq \phi$; and—incidentally—m is self-reproducing provided $m \in O(m)$. This relation, $O(m)$, induces a directed graph on M; by following arrows on this graph we can potentially see what machine types can, directly or indirectly, be constructed by other machine types.
- Let there be a crude ("vague, unscientific and imperfect") measure of complexity on the elements of M—call it $c(m), m \in M$.
- Let $c(m), m \in M$ span a large (ideally, infinite) range, which is to say that there are relatively simple and relatively complex types of machine, and everything in between, included in M.

Now consider the two (highly schematic) graphs shown in Figure 1. In both cases I have shown the putative set M partitioned more or less coarsely by the complexity measure c. That is, the inner rings or subsets are those of small c (low complexity), while the outer, more inclusive rings or subsets include machines of progressively greater c (higher complexity). The graph on the left indicates our "typical" engineering experience: all constructional pathways lead inward (from more complex to less complex). As a result, complexity will always, and unconditionally, degenerate in time. Conversely, the graph on the right indicates the abstract situation posited by our best current theories of biological evolution: at least some edges of the graph (constructional pathways) lead from the inner rings to their outer neighbors. Provided there are sufficient such pathways, then, starting from only the very simplest machines (or organisms), there will be potential constructional pathways whereby even the most complex of machines can eventually be constructed (in time). Thus complexity *might* grow in time.[1]

The problem which this presents is to show that the experience and intuition underlying the graph on the left is mistaken; to show, in other words, how the situation on the right might in fact be realised.

What would it take to solve this problem, even in principle? Well, one would need to exhibit a concrete class of machines M, in sufficient detail to satisfy ourselves that they *are* purely mechanistic; one would need to show that they span a significant range of complexity; and finally, one would have to demonstrate that there *are* constructional pathways leading from the simplest to the most complex (there may, or may not, also be pathways in the other direction, but that is not the point at issue—or at least, not immediately).

I believe that *this* is precisely the problem which von Neumann set out to solve. Furthermore, it seems to me that he did, indeed, solve it; and that it

[1] Note that there are still inward, degenerative, pathways shown here; the claim is only that this graph permits the *possibility* of a growth in complexity; whether it actually *will* grow is a different, and an altogether more difficult, question which I do not attempt to address here.

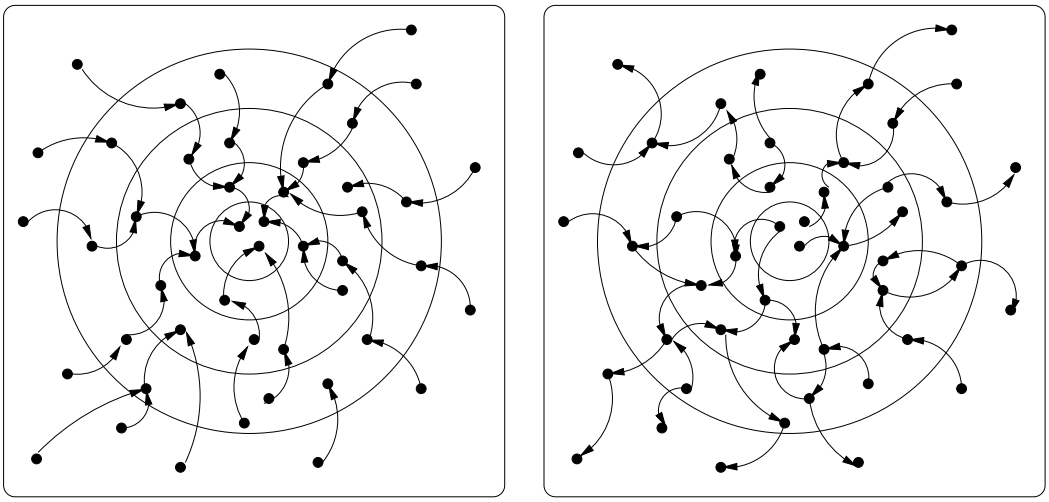

Figure 1: Evolutionary growth of complexity.

is only by seeing his work in *this* light that its "real force" can be understood.

3 Von Neumann's Solution

As to the first claim, that this—rather than "self-reproduction" per se—was von Neumann's problem, it is certainly relevant that he introduced his work in essentially these terms in both his recorded public presentations of it—at the Hixon Symposium in September 1948 (von Neumann, 1951, p. 312) and at the University of Illinois in December 1949 (von Neumann, 1949, pp. 78–79). Granted, in both cases he did *also* refer to the issue of self-reproduction. Specifically, he pointed out that self-reproduction is an example of a constructing machine where the output or offspring is just precisely matched in complexity to the parent—complexity neither increasing nor decreasing. But the critical point here—the significance of self-reproduction—is clearly as a special or "watershed" case of the more general situation: it is interesting precisely, and only, because it may mark a transition from strictly-degenerating to potentially-growing complexity. Conversely, if we encounter a form of "self-reproduction" that is *not* associated with such a transition, then it will not be of relevance to von Neumann's problem at all.

Indeed, von Neumann (1949, p. 86) made this point even more explicitly, when he actually noted the triviality of self-reproduction in "growing crystals" (essentially Burks' 1-state reproducer in a CA). But *pace* Burks, *von Neumann's* resolution of this was not at all to impose a criterion of complexity, but rather to stipulate that the reproductive process should be such as to support "inheritable mutations". Taken in context, I think this must be interpreted as supporting at least some "mutations" where the offspring is of *increased* complexity—which again returns us squarely to the problem of the evolutionary growth of complexity, rather than self-reproduction.

My second claim—that von Neumann actually solved this problem—may seem more far fetched; but let us take it in steps.

First note that while von Neumann exhibited the design and operation of only one machine in detail—his self-reproducing machine—he consistently pointed out how his CA space could support an indefinite variety of quite arbitrary machine configurations and behaviours. In this sense, he certainly exhibited not just a single particular machine, but a whole class of machines (which he identified via the infinite set of arbitrary, finite, "initially quiescent" configurations in his CA). This can correspond to the class M in my problem formulation above (though I will refine this somewhat below).

The next question is whether this class of machines spans a significant range of complexity. Given that we are using only the most vague definition of "complexity" here, the answer can be at best informal and qualitative. To my knowledge, von Neumann did not quite comment on this explicitly. He did explain at length that in setting up this sort of "axiomatic" theory of automata, there is necessarily a degree of arbitrariness in selecting the primi-

tive or atomic parts from which the composite automata will be constructed. And indeed, between 1948 and 1953 he considered a number of alternative approaches before settling on a particular CA formulation (the latter concept having been suggested by Ulam—see Burks, 1966, p. 94). But given the highly complicated (co-ordinated, systematic, purposeful) behaviour of the one machine he did design in detail, it certainly seems to me that if we allow, as von Neumann did, for the construction of machines with indefinitely more parts, in arbitrary configurations, then this set surely *might* span a sufficient range of complexity to meet our requirements.

This then leaves what seems by far the most intractable aspect of the problem: to show that there are constructional pathways leading from the simplest to the most complex machines. At first sight, it seems hopeless to demand a *demonstration* of this; firstly because the measure of complexity is so vague; but secondly because it would seem to demand separate analysis or demonstration of the constructional potentialities for most if not all of the elements of the infinite set M. But this, it seems to me, was where von Neumann's crucial insight occurred.

Drawing on the capability for a single universal Turing machine to carry out the computations of any Turing machine at all, given a suitable description of that machine, von Neumann recognised that a single "general constructive automaton" (von Neumann, 1949, p. 85) might be able to construct "any" machine at all (i.e., any element of his specified M), given a "description tape" of that target machine. This immediately suggests the possibility that certain such special, "programmable", constructor (sub-)systems may enable enormous constructive potentialities—and also, incidentally (?) open up a powerful and *generic* route toward self-reproducing capability.

Von Neumann exploited this general idea by making a minor, but subtle, modification of his general constructive automaton (in a manner incidentally having no analog or significance in pure computation theory): as well as requiring it to *decode* a description to produce an offspring, he required that it must also *copy* the description, attaching this as part of the offspring. As Langton surmised, this combination of decoding and copying potentially has quite dramatic consequences—but the challenge is still to make those consequences explicit.

Let me denote a basic general constructive (decoding and copying) machine by u_0. We require that $u_0 \in M$.[2] Let $d(m), m \in M$ denote a "description of m", (relative to u_0). That is, letting $u_0 \oplus d(m)$ denote the composition of u and a tape describing an arbitrary m, this will result in the construction of an instance of m, via *decoding* of the description, itself composed with a *copy* of the original tape, i.e., $m \oplus d(m)$. We write this as:

$$(u_0 \oplus d(m)) \rightsquigarrow (m \oplus d(m))$$

Since this applies for all $m \in M$ it applies to u_0 itself and we have:

$$(u_0 \oplus d(u_0)) \rightsquigarrow (u_0 \oplus d(u_0))$$

This is the single self-reproducing automaton, which has been commonly identified as von Neumann's "result" or "achievement". Certainly at this point we have a case of construction where complexity has been fully preserved—but the question remains whether this gives us an avenue into the case where complexity can grow.

Now note that u_0 can be combined or augmented with fairly arbitrary ancillary machinery, while still retaining its general constructive ability. That is, we can say that $u_0 \oplus m$ is still a general constructive automaton for "almost all" $m \in M$. The exception would be where m in some sense interferes with or disrupts the normal operation of u_0. Discounting those cases, we can say that the existence of a single general constructive automaton, u_0, actually implies the existence of a whole infinite set of related general constructive automata, $U \subset M$, where $u_m \in U$ has the form $u_m = (u_0 \oplus m)$ (i.e., all the members of this set share the same "general constructive subsystem"). This, in turn, implies the existence of a whole infinite set of self-reproductive automata, S where each $s_m \in S$ has the form $u_m \oplus d(u_m)$. Furthermore, given that U is derived from M just by excluding some exceptional or pathological cases (i.e., where the operation of u_0 would be disrupted), we can say that U, and thus S, will still span essentially the same significant range of complexity as M itself.

Now this is clearly a much stronger result than merely demonstrating the existence of a single self-reproducing machine. In particular, it indicates the possibility of arbitrarily complex machines that are still capable of self-reproduction. This in itself certainly distinguishes von Neumann's result from that of Langton (1984). Although the operation of Langton's machine can be decomposed into copying and

[2] It is neither trivial nor obvious that there exists a general constructive automaton within any given M; this is why the bulk of von Neumann's unfinished manuscript is taken up with the detailed design of a particular example u_0 to establish this result for his particular M.

decoding activities, it does not incorporate anything like a "general constructive automaton". The decoding capability is extremely impoverished so that, in effect, there is only one description that can be effectively processed—that of the one self-reproducing configuration which Langton exhibited.

But in any case, there is still a further critical step in the analysis of von Neumann's work. The goal is not just to show a class of arbitrarily complex constructors capable of constructing machines of equal complexity (which is what self-reproduction illustrates), but to demonstrate constructional pathways whereby complexity can *grow*. And it turns out that this further result can also be demonstrated by the set S.

To see this, we imagine the possibility of perturbations—"mutations"—to the description tape of a machine $s = (u_m \oplus d(u_m)), \in S$. With reasonable assumptions about the description language, it can be arranged that all, or almost all, tape configurations are "legal" (i.e., can be decoded to *some* target machine), and that the tape $d(u_m) = d(u_0 \oplus m)$ can be decomposed into a part, $d(u_0)$, coding for u_0, and a part, $d(m)$, coding for the ancillary machinery, m. That being the case, a more or less arbitrary modification of the $d(m)$ portion of the tape will result in a legal description of some other machine, $d(m')$. The reproductive cycle will then result, not in self-reproduction, but in a "mutant" offspring $s' = (u_{m'} \oplus d(u_{m'}))$; but this is still an element of S, and thus self-reproducing in its own right.[3]

Now these mutations essentially open up additional constructional pathways *between* the elements of S. In particular it is now clear that this can allow incremental, bootstrapping, *increases* in complexity. In effect, the density of these mutational pathways reflects the combinatorics of the description code, so that we can be virtually guaranteed, *without any detailed analysis,* that there will be constructional pathways connecting the entire set S. These will include, of course, degenerating pathways, where the offspring are less complex; but will also include vast numbers of pathways of increasing complexity.

In this way, finally, we see how von Neumann's detailed design of a single machine implies at least a schematic solution of the generic problem of the evolutionary growth of machine complexity.

[3]Indeed, we call such changes "mutations" precisely because they can "breed true" in the offspring.

4 Looking Backwards

While the designation of a distinct field known as *Artificial Life* is comparatively recent (Langton, 1989), I would argue that, looking backwards, it is clear that von Neumann's work properly defines its origin and inspiration. If the formulation above of von Neumann's problem—and his solution to it—is accepted, then a proper assessment of his contribution to the field he effectively founded presents at least three distinct questions:

1. Was von Neumann the *first* one to solve his problem?
2. Has von Neumann presented the *only* known solution to his problem?
3. Is von Neumann's solution the *simplest* known?

My answer to question 1 is clearly an unambiguous "yes"; and it stands as von Neumann's remarkable achievement *both* to have formulated this foundational problem *and* to have actually succeeded in solving it.

Question 2 is slightly more difficult. Given than von Neumann's problem situation has been poorly understood, it follows that subsequent contributors have not generally articulated clearly the relation between their work and this problem.

However, I do not hesitate to say that Thatcher (1970), Codd (1968) and Pesavento (1995) have all offered alternative, or at least significantly refined, solutions. Because these three systems are so very closely related to von Neumann's it follows that they are at least equally satisfactory in solving his problem.

Thatcher provided a somewhat simplified design for the basic general constructive automaton u_0, in the *same* CA space as defined by von Neumann. Codd defined a significantly simpler CA space, which could still accommodate a functionally equivalent general constructive automaton. Much more recently, Pesavento has demonstrated a substantially simplified design for u_0 in a CA space marginally more complicated than von Neumann's; and, further, has actually demonstrated a functioning simulation of this complete u_0.

Conversely, as already indicated, I think we can be clear that Langton's "self-replicating loop" system Langton (1984) does *not* qualify as an alternative solution to von Neumann's. Although it was derived from Codd's earlier system, and although he exhibits a self-reproducing configuration (involving "copying" and "decoding") this does not embody

anything like a "general constructive automaton" and therefore has little or no evolutionary potential.

Another possible candidate solution to von Neumann's problem might be that of Berlekamp et al. (1982). There it is claimed that a general constructive automaton, of comparable functionality to von Neumann's u_0, can be realized in Conway's well known "Game of Life" (GOL) CA. This arguably takes Codd's (and thus von Neumann's) result to some kind of limiting case, as the GOL cells have the absolute minimum of only 2 states. On the other hand, in contrast to the fully detailed designs for u_0 presented by Thatcher, Codd and Pesavento, Berlekamp et. al. provide only a very sketchy outline to indicate the *possibility* of an equivalent automaton in GOL. Moreover, because GOL is so radically simplified (as a CA) it would not be quite trivial to establish that it supports embedded automata of comparable range of "complexity" to those of Von Neumann etc.

This last issue affects the comparative evaluation of certain other systems even more seriously. I have in mind particularly `Tierra` (Ray, 1992) and related systems.

`Tierra` is, roughly, a shared memory parallel computer, with fairly simple "von Neumann style" processors.[4] Ray has exhibited processes ("machines" or "automata" in von Neumann's sense) in this framework which are capable of self-reproduction; and which, moreover, exhibit clear and significant evolutionary change. However, because the `Tierra` system is so very different in style from von Neumann's (or, indeed, *any* CA of 2 or more dimensions), it is very difficult to make even informal and qualitative comparisons of automata "complexity" between these systems.

I also have another quite different, and perhaps more important, misgiving about whether `Tierra` should be regarded as offering an equally satisfactory alternative solution to von Neumann's problem.

I originally expressed this reservation by stating that `Tierra` uses "... a form of self-reproduction based on self-inspection (rather than a properly genetic system in the von Neumann sense)" (McMullin, 1992, Chapter 4). The implication was that the self-reproducing entities in `Tierra` lack the important distinction between "copying" and "decoding"—and that this perhaps affects evolutionary potential in the system. However Taylor (1999) has since persuaded me that this way of presenting the matter is, at best, unclear.

It might be more precise to say that `Tierra` *does* incorporate a distinction between copying and decoding processes: but that the "decoding" is represented by the execution of instructions by the `Tierra` processors. Thus, the decoding is hardwired, whereas, in von Neumann's system, the decoding is itself a product of the specific design of u_0, and, as such, is at least *potentially* mutable. In this way, von Neumann's system allows for what I have called *Genetic Relativism* (McMullin, 1992, Chapter 4), where the actual "decoding" map can itself evolve. It seems to me that this allows for a more profound form or degree of evolvability in von Neumann's system compared to `Tierra`. Having said that, I should also note that von Neumann himself seems to have discounted this idea. He stated explicitly that mutations affecting that part of a descriptor coding for u_0 would result in the production of "sterile" offspring (von Neumann, 1949, p. 86)—and would thus have no evolutionary potential at all. Clearly, on this specific point, I disagree with von Neumann, and consider that he actually underestimated the force of his own design in this particular respect.

As to my question 3 above—whether von Neumann's solution is the "simplest"—this clearly does not admit of a clearcut answer, given the informality of our notions of simplicity and complexity. Certainly, the alternative solutions of Thatcher, Codd and Pesavento can all be said to be *somewhat* simpler in at least some respects. Similarly, if the outline design of Berlekamp et. al. is accepted then it is *much* simpler in one particular dimension (the number of cell states). Nonetheless, in considering the overall development of the field I would say that none of these "simplifications" either undermine, or significantly extend, von Neumann's seminal results.

5 Conclusion: Looking Forwards

By re-examining von Neumann's work in the light of his own description of the problem he was working on, I have tried to show that there is much more substance to it than has been generally recognised. However, as Popper (1976) has emphasised, the scientific enterprise is iterative: the solution of any given problem gives rise to a new problem situation, which is to say new problems to be addressed. It seems to me that von Neumann's work is particularly fruitful in this regard, as it poses a number of new and profound questions which need to be ad-

[4] Here, of course, I am referring to the so-called "von Neumann computer architecture" (von Neumann, 1945)—which was not, by any means, solely von Neumann's invention—rather than to his very different work on cellular automata.

dressed by the (still) fledgling field of Artificial Life. Among these are:

1. Can we clarify (or even formalise) the notion of complexity. This of course is not specifically a problem for Artificial Life, but rather underlies the entire discipline of evolutionary biology (Maynard Smith, 1969).

2. What is the precise significance of the self-reproductive capability of von Neumann's machines? Technically, based on the idea of a "general constructive automaton", growth of complexity *per se* could take place in isolation from self-reproduction. One simple scenario here would be to imagine u_0 being made to simply grind through the (countable) infinity of all description tapes, constructing every described automaton in turn. This would require only a trivial extension of the capabilities of u_0. While this would fail in practice due to the essential fragility of von Neumann's automata (discussed further below) it is not clear whether there is any *fundamental* problem with this general idea. Nonetheless, it seems clear that if the growth of complexity is to involve *Darwinian* evolution, then self-reproduction surely is a necessary additional requirement.

3. The problem of identity or individuality. In the alife systems discussion above, the question of what constitutes a distinct individual "machine" or "automaton" is addressed in a completely ad hoc manner. In a real sense, these putative individuals exist as such only in the eye of the human observer. Whereas, the notion of a self-defining identity, which demarcates itself from its ambiance, is arguably essential to the very idea of a biological organism. Some early and provocative work on this question, including simulation models overtly reminiscent of von Neumann's CA, was carried out by Varela et al. (1974). However, there has been relatively little further development of this line since.

4. The evolutionary boot-strapping problem: in the von Neumann framework, at least, u_0 is already a very complicated entity. It certainly *seems* implausible that it could occur spontaneously or by chance. Similarly, in real biology, the modern (self-consistent!) genetic system could not have plausibly arisen by chance (Cairns-Smith, 1982). It seems that we must therefore assume that something like u_0 (or a full blown genetic system) must itself be the product of an extended evolutionary process. Of course, the problem with this—and a major part of von Neumann's own result—is that it seems that something like a genetic system is a *pre-requisite* to any such evolutionary process.

5. Perhaps most importantly of all, what further conditions are required to enable an *actual*, as opposed to merely *potential* growth of complexity? It was well known even to von Neumann himself that his system would not *in practice* exhibit any evolutionary growth of complexity. The proximate reason is that, in his CA framework, all automata of any significant scale are extremely *fragile*: that is, they are very easily disrupted even by minimal perturbation of the external environment. The upshot is after the completion of even a single cycle of self-reproduction the parent and offspring would almost immediately perturb, and thus effectively destroy, each other. This can be avoided by ad hoc mechanisms to prevent all interaction (again, this was suggested by von Neumann, and since shown in practice by Langton and others). However, eliminating interaction eliminates the grist from the darwinian mill, and is thus a non-solution to the substantive problem (growth of complexity). `Tierra` offers a somewhat different approach, whereby the integrity of the automata (process images) is "protected" by the underlying operating system, while still allowing some forms of significant interaction. This has been very fruitful in allowing the demonstration of *some* significant evolutionary phenomena. However, any serious claim to substantively model real biological organisms will inevitably have to confront their capacity for *self* maintenance and repair in the face of continuous perturbation and material exchange with their environments. This, in turn, is clearly also related to the earlier problem of biological individuality.

In conclusion, it seems to me that von Neumann's work in Artificial Life—properly understood—is as profound and important today as it was half a century ago; and that it should continue to provide structure, insight and inspiration to the field for many years to come.

References

Berlekamp, E. R., Conway, J. H. and Guy, R. K. (1982), What is Life?, *in* 'Winning Ways for your Mathematical Plays', Vol. 2, Academic Press, London, chapter 25, pp. 817–850.

Burks, A. W. (1970*a*), Von Neumann's Self-Reproducing Automata, in *Essays on Cellular Automata* Burks (1970*b*), pp. 3–64 (Essay One).

Burks, A. W., ed. (1966), *Theory of Self-Reproducing Automata [by] John von Neumann*, University of Illinois Press, Urbana.

Burks, A. W., ed. (1970*b*), *Essays on Cellular Automata*, University of Illinois Press, Urbana.

Cairns-Smith, A. G. (1982), *Genetic Takeover and the Mineral Origins of Life*, Cambridge University Press, Cambridge.

Codd, E. F. (1968), *Cellular Automata*, ACM Monograph Series, Academic Press, Inc., New York.

Herman, G. T. (1973), 'On Universal Computer-Constructors', *Information Processing Letters* **2**, 61–64.

Langton, C. G. (1984), 'Self-Reproduction in Cellular Automata', *Physica* **10D**, 135–144.

Langton, C. G. (1989), Artificial Life, *in* C. G. Langton, ed., 'Artifical Life', Vol. VI of *Series: Sante Fe Institute Studies in the Sciences of Complexity*, Addison-Wesley Publishing Company, Inc., Redwood City, California, pp. 1–47.

Maynard Smith, J. (1969), The Status of Neo-Darwinism, *in* C. H. Waddington, ed., 'Towards a Theoretical Biology, 2: Sketches', Edinburgh University Press, Edinburgh, pp. 82–89. This paper is also accompanied by various addenda and comments (pages 90–105 of the same work).

McMullin, F. B. V. (1992), Artificial Knowledge: An Evolutionary Approach, PhD thesis, Ollscoil na hÉireann, The National University of Ireland, University College Dublin, Department of Computer Science.
`http://www.eeng.dcu.ie/~alife/bmcm_phd/`

Minsky, M. L. (1967), *Computation: Finite and Infinite Machines*, Prentice-Hall Series in Automatic Computation, Prentice-Hall Inc., Englewood Cliffs, New Jersey.

Pesavento, U. (1995), 'An implementation of von Neumann's self-reproducing machine', *Artificial Life* **2**(4), 337–354.

Popper, K. R. (1976), *Unended Quest*, Fontana/William Collins Sons & Co. Ltd, Glasgow.

Ray, T. S. (1992), An approach to the synthesis of life, *in* C. G. Langton, C. Taylor, J. D. Farmer and S. Rasmussen, eds, 'Artifical Life II', Vol. X of *Series: Sante Fe Institute Studies in the Sciences of Complexity*, Addison-Wesley Publishing Company, Inc., Redwood City, California, pp. 371–408. Proceedings of the workshop on Artificial Life held February, 1990, in Sante Fe, New Mexico.

Taub, A. H., ed. (1961), *John von Neumann: Collected Works. Volume V: Design of Computers, Theory of Automata and Numerical Analysis*, Pergamon Press, Oxford.

Taylor, T. J. (1999), From Artificial Evolution to Artificial Life, PhD thesis, University of Edinburgh.
`http://www.dai.ed.ac.uk/daidb/homes/timt/papers/thesis/html/main.html`

Thatcher, J. W. (1970), Universality in the von Neumann Cellular Model, *in* Burks (1970*b*), pp. 132–186 (Essay Five).

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. (1945), First Draft of a Report on the EDVAC. The (corrected) version at the URL below has been formally published in the IEEE Annals of the History of Computing, **15**(4), 1993.
`http://qss.stanford.edu/~godfrey/vonNeumann/vnedvac.pdf`

von Neumann, J. (1949), Theory and Organization of Complicated Automata, *in* Burks (1966), pp. 29–87 (Part One). Based on transcripts of lectures delivered at the University of Illinois, in December 1949. Edited for publication by A.W. Burks.

von Neumann, J. (1951), The General and Logical Theory of Automata, *in* Taub (1961), chapter 9, pp. 288–328. Delivered at the Hixon Symposium, September 1948; first published 1951 as *pages 1–41 of:* L. Jeffress, A. (ed), *Cerebral Mechanisms in Behavior*, New York: John Wiley.

Acknowledgements

Many of the ideas presented here were first explored in my PhD thesis (McMullin, 1992); I was privileged to carry out that work under the supervision of the late John Kelly. A large number of people have since helped and challenged me in trying to refine the analysis and arguments. I am particularly indebted to Chris Langton, Glen Ropella, Noel Murphy, Tim Taylor, Mark Bedau and Moshe Sipper; the latter also provided critical encouragement to write this particular paper. I am grateful also to the ALife VII reviewers for comprehensive and constructive criticism. Financial support for the work has been provided by the Research Institute in Networks and Communications Engineering (RINCE) at Dublin City University.

Revision History

First published in August 2000 as Technical Report No. `bmcm-2000-01` of the Artificial Life Laboratory of Dublin City University. First presented at the conference *Artificial Life VII* August 1–6, 2000, Reed College, Portland, Oregon. Essentially the same text appears in *Artificial Life VII: Proceedings of the Seventh International Conference,* edited by Mark A. Bedau, John S. McCaskill, Norman H. Packard, and Steen Rasmussen (MIT Press, 2000), pp. 467–476.

Author Contact Information

Barry McMullin

DCU Alife Laboratory
Research Institute for Networks and Communications Engineering (RINCE)
Dublin City University
Dublin 9
Ireland.

Voice:	`+353-1-700-5432`
Fax:	`+353-1-700-5508`
E-mail:	`Barry.McMullin@dcu.ie`
Web:	`http://www.eeng.dcu.ie/~mcmullin`

Online Retrieval

The resources comprising this paper are retrievable in various formats via:

- `http://www.eeng.dcu.ie/~alife/bmcm-2000-01/`

Related Online Resources

- Presentation Slides:
  - `http://www.eeng.dcu.ie/~alife/talks/alife7/vn-complexity/`
- Reviewer Comments on draft manuscript:
  - `http://www.eeng.dcu.ie/~alife/bmcm-2000-01/review-246.html`
  - `http://www.eeng.dcu.ie/~alife/bmcm-2000-01/review-250.html`
- ALife7 Evolvability Workshop Paper: *The Von Neumann Self-reproducing Architecture, Genetic Relativism and Evolvability:*
  - `http://www.eeng.dcu.ie/~alife/bmcm-2000-02/`
- DCU Alife Laboratory:
  - `http://www.eeng.dcu.ie/~alife/`
- Research Institute for Networks and Communications Engineering (RINCE):
  - `http://www.rince.ie/`

Copyright

This report is copyright ©2000 by Barry McMullin. Certain rights have been assigned to The MIT Press in connection with its copyright in the edited volume *Artificial Life VII: Proceedings of the Seventh International Conference,* edited by Mark A. Bedau, John S. McCaskill, Norman H. Packard, and Steen Rasmussen (MIT Press, 2000).

Permission is hereby granted to private individuals to access, copy and distribute this work, for purposes of private study only, provided that the distribution is complete and unmodified, is accompanied by this copyright notice, and that no charges are levied. The work may *not* be accessed or copied, in whole or in part, for commercial purposes, except with the prior written permission of the author.

All other rights reserved.