

Artificial Life Meets Computational Creativity?*

Presented at
Computational Creativity : An Interdisciplinary Approach[†]

12–17 July, 2009, Dagstuhl, Germany

Barry McMullin[‡]

Dublin City University
July 2009

Abstract

I review the history of work in Artificial Life on the problem of the open-ended evolutionary growth of complexity in computational worlds. This is then put into the context of evolutionary epistemology and human creativity.

*Parts of this text are based on a book chapter, co-authored with Wolfgang Banzhaf, for the forthcoming Springer *Handbook of Natural Computing* (Rozenberg et al., 2010).

[†]<http://www.dagstuhl.de/en/program/calendar/semhp/?semnr=09291>

[‡]<http://www.eeng.dcu.ie/~mcmullin/>

1 Introduction

By extending the horizons of empirical research in biology beyond the territory currently circumscribed by life-as-we-know-it, the study of Artificial Life gives us access to the domain of life-as-it-could-be, and it is within this vastly larger domain that we must ground general theories of biology and in which we will discover practical and useful applications of biology in our engineering endeavors.

—Chris Langton¹

The domain of this Dagstuhl seminar has been the use or application of computers in realising “creativity” — understood particularly as what is recognised as creative in human activity. This may take the form of computer aids or amplifiers of human creativity; but also standalone computational systems which are claimed to demonstrate some autonomous creative phenomena. In this contribution I focus particularly on the latter case, in the form of autonomous, computationally realised, “virtual worlds”. I review attempts to understand, through building such systems, the conditions and constraints on such spontaneous creativity. This work is typically classified as falling within the field of *Artificial Life*.

2 Constructive and Autopoietic Systems

The notion of creativity is often related to more general concepts of emergence and complex organisation. Within the field of Artificial Life, this bears on what may be called “self-constructing” systems. Natural life, at all hierarchical levels, exhibits a characteristic ability to maintain its system organisation while simultaneously turning over components at the lower hierarchical levels. Thus, a cell can stably maintain itself (and even grow and divide) while continuously regenerating all its significant molecular components (metabolising environmental substrates as necessary); similarly, a multi-cellular organism maintains its organisation while constantly replacing its component cells; and a social insect colony may long outlive most, if not all, of its constituent individual organisms. While this is a typical phenomenon of natural life, it is very different from the behaviour of conventional engineered systems. The latter may well have complex hierarchical organisation, but this organisation is typically static: the system is assembled once, and then retains its fixed compositional structure for its functional lifetime. Indeed, it is typical of engineered systems that failure of any component, at any hierarchical level, will at least significantly impair system level function, and will commonly result in full system failure. Although so-called “fault tolerant” systems may be engineered to incorporate significant redundancy, this often involves static assembly of additional components at manufacture time. Functional lifetime is extended as failing components are automatically removed from service; but once this pre-built redundancy is exhausted, system failure follows from any further

¹<http://www.biota.org/papers/cglalife.html>

component failure. Further, even such redundant systems are generally vulnerable to failures in the “failure detecting” components which cannot also simply be duplicated without potentially entering into infinite regress.

This is a complex problem to address, but natural life provides a “proof of principle” that such self-constructing organisations are possible. Artificial Life certainly offers a potentially very fruitful avenue for further investigation. In particular, Artificial Life allows “virtual worlds” to be formulated in which problems of self-construction can be simplified, and the core organisational mechanisms can be exposed and subjected both to mathematical analysis and experimental exploration. While a variety of work in Artificial Life bears on this, some of the most direct contributions might be summarised as follows:

- The Chilean biologists Maturana and Varela pioneered the abstract concept of *autopoiesis* (literally “self-production”) as a description of the core mechanism of organisational self-maintenance in biological cells (Varela et al., 1974). Almost 15 years before the modern computational study of Artificial Life was even named (Langton, 1989), they were already using a molecular level, agent-based, abstract chemistry to give a concrete demonstration of this abstract theory of cellular organisation. In essence, the proposal is that biological cells are dynamic, self-sustaining, chemical networks (an abstraction of cell metabolism), which also create and maintain a spatial boundary (an abstraction of a cell membrane), where the contained network and the boundary reciprocally rely on each other for stability. While some technical deficiencies were identified in the original presentation of this work (McMullin and Varela, 1997), it has given rise to a sustained and continuing programme of active research (McMullin, 2004).
- Also in the 1970’s, Holland proposed the so-called α -universes (Holland, 1976) and, independently, Hofstadter described the *Typogenetics* system (Hofstadter, 1979). Both involved one-dimensional fragments of computer code, which could interact with each other. In both cases, they were explicitly inspired by aspects of the molecular replication and translation machinery of biology, and were concerned with understanding the reflexive interactions that arise when the “same” class of entities (molecules) can sometimes function as executable “code” (enzyme/protein) and sometimes as literal “data” (nucleic acid), and the possibility of collective self-reproduction and maintenance. These were, however, restricted to theoretical/analytic treatments, and were not subject to empirical investigation until much later (McMullin, 1992; Snare, 1999; Kvasnicka et al., 2001).
- This same tension between syntax and semantics in materially instantiated dynamic systems is also at the core of Pattee’s analysis of what he terms “semantic closure” (Rocha, 2000; Cariani, 1992).
- A related, but independent line of investigation was pioneered by Kauffman and others, in the form of “collectively autocatalytic sets” (Farmer et al., 1986; Kauffman, 1993). Such a set is formally similar to (though more strictly defined than)

the contained chemical network of an autopoietic system, but without the requirement for a self-generated spatial boundary. The key result here was the demonstration, in completely abstract virtual “chemistries” that such collectively autocatalytic sets can spontaneously arise, quickly and under relatively weak constraints on the underlying chemistry.

- Fontana and Buss self-consciously launched a mathematical and computational investigation of “constructive dynamical systems”, based around the so-called *Alchemy* system (Fontana and Buss, 1994). These systems deliberately diverge from the classical concept of “dynamical systems” by relaxing the normally strict demarcation between “state variable” and “dynamic law”. The authors demonstrated and analysed a number of organisational phenomena, including forms of “self-sustaining closure” where more-or-less complex dynamic aggregates of components successfully sustain themselves, even as all the individual components are continuously diluted or degraded.
- As a comprehensive overview treatment, Dittrich et al. (2001) provide a summary of the general approach of building computational models of arbitrary “artificial chemistries”, as a platform for investigating these and other problems. More recently, Dittrich and Speroni have developed a primarily algebraic mathematical analysis in a comprehensive way, in the form of *Chemical Organisation Theory* (Dittrich and di Fenizio, 2007). However, while an important advance, the difficulties of incorporating dynamics (chemical kinetics) and the interaction of chemical self-production and spatial demarcation (the self-constructed “boundary” of autopoietic theory) remain substantial.
- A quite different and more radical approach to the problem of biological self-maintenance was proposed by the theoretical biologist Rosen over an extended series of works (Rosen, 1959, 1972, 1985, 1991). In brief, Rosen argues that the self-constructing organisation characteristic of natural living systems which he termed “closure under efficient causation” transcends the possibilities of conventional dynamical systems in a fundamental way; with the consequence that it cannot be realised in any “computational” universe. This is a complex and contentious theoretical claim which, not surprisingly, continues to be the focus of considerable critique and criticism (e.g., Chu and Ho, 2006). Rosen’s work was also a strong influence in Kampis’ development of a comprehensive mathematical treatment of what he terms “self-modifying” systems (Kampis, 1991).

3 A Core Problem of Creation: Computing the Evolutionary Growth of Complexity

There is a concept which will be quite useful here, of which we have a certain intuitive idea, but which is vague, unscientific, and imperfect ... I know no adequate name for it, but it is best described by calling it "complication." It is effectivity in complication, or the potentiality to do things. 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.

— John von Neumann (von Neumann, 1949, p. 78, emphasis added)

Arguably the most fundamental problem in the theory of biology is that of the growth of "complexity"; with "complexity" understood in the sense explained by von Neumann above, of the ability to do complicated things. While this is certainly an informal definition, it suffices to clearly demarcate this notion from purely syntactic, combinatorial or computational concepts of complexity. Complexity in this sense is also clearly related to at least some meanings of "creativity".

Von Neumann is a key figure in the early development of abstract computational modelling approaches to understanding biological phenomena; indeed, in this respect he can be considered as having instigated the earliest investigations in Artificial Life, as the term is now understood (McMullin, 2000a). In particular, in the late 1940's he started to draw attention to an apparent paradox arising from the contrast between any mechanistic (i.e., not vitalistic) understanding of living organisms and common experience of engineering artificial mechanisms or automata (von Neumann, 1949). While it is quite generally possible to design machines that construct other machines, this process is normally *degenerative in complexity*: a machine of a given complexity (such as an automated factory) can only construct machines of comparatively lower complexity (such as cars or phones or televisions etc.). Whereas, if the theory of Darwinian evolution is accepted, then biological "machines" (organisms) must be capable of constructing descendant machines of greater complexity. Granted, these increases in complexity may occur only in very small increments, and even then only in some lineages and accumulating over geological periods of time; but nonetheless it must be possible, in principle, for biological organisms to give rise to offspring more complex than themselves.

Von Neumann quickly developed an outline resolution of this paradox, in the abstract form of what he termed a "general constructive automaton". This was inspired by Turing's earlier formulation of a universal computing automaton (Turing, 1936), and also by von Neumann's own contemporaneous involvement in designing and building the earliest electronic stored program digital computers (von Neumann, 1945). In essence, a general constructive automaton is a *programmable* constructor or assembler, capable of constructing an indefinitely large set of target automata — loosely, a "universal" set comprised of all automata that could be assembled from a given, finite, set of primitive components. Significantly, von Neumann hypothesised that a

general constructive automaton, relative to a specific set of components, might itself be realisable as an assembly of these same components. With this conceptual architecture and some relatively minor technical elaboration, von Neumann showed that this then would give rise to an indefinitely large set of automata, spanning an indefinite large range of complexity, all of which would be fully connected by a network of heritable “mutations” (essentially, chance variations in the construction programs). That is, there would exist mutational pathways leading from the simplest to arbitrarily complex members of this set; with all these machines also, incidentally, being capable of self-reproduction.²

In retrospect this was already an astonishing achievement. Von Neumann effectively described the abstract architecture of biological self-reproduction, based on separate processes of syntactic copying (replication) and semantic “decoding” (translation) of a passive information carrier; he showed how this architecture supported self-reproduction, heritable mutation, and thus evolutionary growth of complexity; and he presented all this, at least in outline form, in 1948, five years before the chemical structure of DNA was even identified (von Neumann, 1949).

However, as yet this was still only a sketch of a solution. To make it properly convincing, he needed to present a concrete set of “primitive parts” and show that with these, it would actually be possible to realise an example of a general constructive automaton. This proved to be a complicated and potentially intractable problem: any plausible “real world” set of components would introduce many ancillary complications — mechanics, thermodynamics, etc. Following a suggestion from Ulam, von Neumann instead formulated an artificial, virtual, world that would simplify away these ancillary complications. He proposed a two dimensional “tessellation automaton”, or, as it is now called, a *cellular automaton*. Within this simplified virtual world, von Neumann then successfully developed a fully detailed design for an example of a general constructive automaton. He had planned this as a first step in a much more general and comprehensive “theory of automata”. However, he put his unfinished manuscript aside in 1953; and due to his untimely death in 1957, was never to return to this project. The manuscript was subsequently edited by Burks and formally published in 1966 (Burks, 1966).

Von Neumann’s achievement, both in theoretical insight and technical implementation, was considerable, and it gave rise to an extended programme of further research. Many different cellular automata worlds have been formulated and investigated, and many variants on von Neumann’s architecture for evolvable self-reproducing automata have been demonstrated (Thatcher, 1970; Codd, 1968; Langton, 1984). Conway’s *Game of Life* has particularly popularised the study of “life-life” phenomena in cellular automata worlds (Berlekamp et al., 1982). More recently, it has become feasible to actually implement von Neumann’s original designs on real computers (Pesavento, 1995). Cellular automata have also, of course, featured as inspiration and realisation of various art works and installations (e.g., Brown, 1996; Satomi and Sommerer, 2007).

²While terminology in the field is not completely consistent, in this paper I will reserve “(self-)reproduction” to mean this von Neumann style process, involving separate actions of copying and decoding; whereas “(self-)replication” is used to denote a process of copying only.

Nonetheless, seminal as this work was, it also explicitly left certain questions still very much unresolved. In particular, as von Neumann himself was clearly aware, the mere existence of mutational pathways from simple to complex automata does not mean that evolution will actually follow such pathways. Indeed, in von Neumann's original cellular automaton world, the embedded self-reproducing machine was extremely fragile. It could only operate successfully at all on the condition that the rest of the world was completely quiescent; and, of course, once it completes one round of reproduction, that could no longer be the case. Parent and offspring would then interfere with each other's structure and operation, and both would rapidly disintegrate. While this problem could be superficially avoided by constraining each complete, successively constructed automaton to restrict its further operation to separate regions of the space, it rules out any possibility of natural selection; and, in any case, cannot be sustained in the long term if there is any possibility of stochastic, component level, malfunction or failure (which itself is actually *desirable*, as a source of the random variation that is the grist to the mill of Darwinian evolution).

4 Coreworlds: Spontaneous Evolution of Computer Programs

Nothing in Biology makes sense except in the light of evolution.

— *Theodosius Dobzhansky* 1973

Discovering how to make such self-reproducing patterns more robust so that they can evolve to increasingly more complex states is probably the central problem in the study of artificial life.

— *J. Dooyne Farmer and Alletta d'A. Belin* 1992

Although the cellular automata models of von Neumann and his successors have not, to date, demonstrated extended evolutionary dynamics, various other Artificial life models have addressed evolution more directly, through attempts to demonstrate evolution among populations of virtual (software) agents in virtual worlds. The most systematically investigated framework is to envisage these as small machine code programs, each executed by a separate, parallel, processor, but all sharing a single main memory ("core"). Generically, these approaches can be referred to as *coreworlds*. In principle, the parallelism could be implemented directly with a sufficiently large pool of hardware processing elements, dynamically allocated as agents are created or destroyed; but it is typically just realised with conventional timeslice-based multithreading of a much smaller pool of hardware CPUs (often just one).

The earliest work adopting this methodological approach can be traced to the seminal investigations of Barricelli (1957, 1963) from 1953-1957. In fact, this work was carried out on the original computer designed and built at the Princeton Institute for Advanced Studies (IAS) between 1946 and 1952 under the direction of von Neumann. Barricelli was a visitor to the institute at von Neumann's invitation; and took the opportunity to conduct an investigation of the role of "symbiogenesis" in the origin of

life. He programmed the IAS computer to directly model patterns of numbers which could interact with each other within a circular array (effectively the main memory system of the machine), according to some fixed local rules. This is conceptually similar to a one-dimensional cellular automaton. Although these patterns of numbers did not have the general computing ability of Turing machines, Barricelli demonstrated the existence of patterns which could self-replicate, as well as various forms of competitive and co-operative (symbiotic) interaction between patterns. Somewhat strangely, in his published description of this work, Barricelli only obliquely related it to von Neumann's own contemporary work on self-reproducing and evolving automata. In any case, the insight and innovation of this work was not widely recognised at the time, and has only recently been properly re-instated through George Dyson's investigations of the early history of digital computing at the IAS (Dyson, 1997, 1998).

The first attempt to study the dynamics of competition among co-resident self-replicating *programs* in a shared memory was pioneered by Vyssotsky et. al. at Bell Labs in the early 1960's (though not published until 1972), and appears to have been independent of Barricelli's study (Vyssotsky, 1972). This was effectively created as a form of computer programming game in which different programmers provided hand-coded programs, which were then instantiated in a common core, and executed to see which would survive longest. Again, self-replication was a key feature. An interesting aspect was that, unlike almost all later work, these programs were implemented in the native (IBM 7090) machine code, rather than in a virtual machine code to be executed by an interpreter. This allowed comparatively much higher execution speed, but with some associated limitations. For example, as programs could not be written to be position-independent, successful self-replication required active relocation processing, as opposed to simple self-copying. In this early implementation, a "lethal" program was relatively quickly developed by R.H. Morris, and the original game was put aside.

We should mention again here also the α -universes described by Holland in the mid-1970's (Holland, 1976). Although formulated independently, and with quite different motivations, this proposal has similarities to both the Barricelli and Vyssotsky systems. In particular, it again envisaged an essentially linear core memory inhabited by concurrently executing fragments of computer code. However, in other ways this was actually much closer to Barricelli's work in that the specific intention was to investigate the spontaneous emergence of crude collective self-reproduction activity, rather than simple competition between programs pre-coded by human programmers; but, in any case, this system was not experimentally investigated until much later (McMullin, 1992).

The more direct descendent of the Vyssotsky system was the *Core War* game, developed by Dewdney and others in the early 1980's. This now relied on an interpreter and offered much more varied gameplay opportunities (Dewdney, 1984). Following the establishment of an international tournament (Dewdney, 1987), *Core War* has had a sustained following, and remains active to this day.³

Although the *Core War* framework is predicated on the idea of human programmers coding the competing programs, it does also naturally lead to the question of

³<http://corewars.org/>

whether, under conditions of potential mutation during self-replication, there could be an autonomous and sustained evolutionary process, not relying on programmer intervention. This was explicitly studied by Rasmussen et. al. in a system they called the *Coreworld* (Rasmussen et al., 1990). In this instantiation however, the system suffered from a similar limitation to that of the von Neumann style cellular automata systems. While the world could be seeded with an initial self-replicating program, as its offspring filled up the core they quickly begin overwriting each other and the self-replicating functionality and thus the potential for sustained Darwinian evolution was lost. This issue was addressed by Ray, in the *Tierra* system, developed and investigated in the early 1990's (Ray, 1992). *Tierra* was based on the idea of competing and interacting self-replicating programs in a shared core memory. Programs were now given an ability to employ "memory protection" to limit overwriting by other programs. On its own, this would simply lead to the core filling with programs once, and then no further self-replication (much less evolution) would be possible. To overcome this, Ray added a mechanism for random removal of programs (the so-called "reaper"). This addition allowed for continuing self-replication and the possibility for longer term evolution. He also adopted a significantly different instruction set, incorporating a biologically inspired "template" addressing mode. With these innovations, Ray was able to demonstrate an extended process of evolution, with multiple episodes of Darwinian natural selection. The system produced a surprising array of evolutionary phenomena, including optimisation of program length, obligate parasitism and facultative "hyper-parasitism" (as a defence against obligate parasites).

Tierra has led to a wide variety of related work, such as:

- The *Avida* system of Adami and Brown which introduced a more conventional spatial 2-D spatial world in which programs would colonise nodes, and also be evaluated for their ability to complete extrinsically defined computing tasks (Adami and Brown, 1994).
- A proposed extension to *Tierra* into a multiprocessor networked environment ultimately to be distributed across the wider Internet, under suitable "sandbox" execution controls (Ray, 1995).
- *Nanopond*⁴, a highly simplified implementation of a 2-D spatial "program evolution" system, drawing on elements of both *Tierra* and *Avida*, but implemented in less than 1,000 lines of C source code.
- The *Amoeba* system, developed by Pargellis, which specifically demonstrated the possibility of spontaneous emergence of a self-replicating program in coreworld-like systems (Pargellis, 2001).

⁴<http://adam.ierymenko.name/nanopond.shtml>

Evolution in coreworlds can also usefully be compared and contrasted with several other related but divergent research fields and investigations:

- There has been a separate line of line of research into the logic of so-called *Quine* programs. These are defined as computer programs which produce their own source code as output. The original discussion is generally credited to Bratley and Millo (1972), though the term “Quine program” is usually attributed to Hofstadter (1979). The existence of Quine programs in general purpose programming languages essentially follows from Kleene’s recursion theorem. But while the formulation of specific Quine programs in different languages has been a hobby among programmers, these programs have not generally been embedded within an execution environment in which exponential growth, Darwinian selection, and evolution could actually take place. It has not therefore, particularly influenced research approaches in Artificial Life. That said, through its definitional restriction to programs that must self-reproduce without having direct access to self-inspection, the study of Quine programs draws attention to the logical distinction between properly von Neumann style genetic “self-reproduction” and the “self-replication” by copying normally adopted in coreworlds. It is, in fact, a very open question as to what the effect of this architectural choice is on evolutionary potential (McMullin et al., 2001).
- The replication and propagation of programs in controlled coreworld environments obviously also leads to the idea of self-replicating programs propagating in open networked computer environments. This is the realm of computer malware — computer viruses, worms etc. And indeed, this connection was explicitly made by Dewdney already in the immediate aftermath of the Morris Internet worm incident in 1988 (Dewdney, 1989). However, while malware development certainly involves an arms race between the human developers on both sides, there is to date no evidence of effective *autonomous* evolution of “free-living” malware.
- There is significant overlap in inspiration between evolution in coreworld-like systems, and work in evolutionary algorithms, and, especially, *genetic programming* (GP) (Koza, 1992; Willis et al., 1997; Banzhaf et al., 1998). This, in turn, has a long history extending back at least to Samuel’s late 1950’s investigations in machine learning (Arthur, 1959). However, the major distinction is that GP is generally driven, directly or indirectly, by an externally provided evaluation function (used as an extrinsic “fitness” to drive an imposed evolutionary algorithm) rather than the coreworld approach which is to investigate spontaneous autonomous evolutionary dynamics, in which the software agents are responsible for their own replication and relative fitness emerges from their bottom-up ecological interactions.
- Investigations into the evolution of co-operation by Axelrod and others should also be mentioned here (Axelrod, 1984, 1987). Although not generally using a

coreworld-like framework, and focussed on a relatively narrowly defined pattern of interaction (the iterated prisoner's dilemma), using an extrinsically applied genetic algorithm, this examination of the problem of co-operation, and co-evolution between "co-operating" and "defecting" strategies, has been very influential in the wider fields of political science, ecology, and complexity theory.

To summarise, the general experience of the investigation to date of evolution in coreworld-like systems is that the evolutionary potential of these systems is interesting *but still strictly limited*. Indeed, they can be viewed as formally similar to pure artificial "replicator chemistries", comparable to the systems already mentioned in section 2 above, and exhibiting similar phenomena of collective autocatalytic closure and self-maintenance (McMullin, 2000b). Thus, both von Neumann style genetic self-reproduction in cellular automaton worlds, and replication by simple copying or self-inspection in coreworlds, naturally lead directly into the problems of biological robustness, self-maintenance, and hierarchical organisation already discussed above. The integration of self-reproducing or self-replicating programs with self-maintenance and individuation (autopoiesis), and the demonstration of sustained evolutionary growth of complexity in a purely virtual world remains perhaps *the key "grand challenge"* problem in the field of Artificial Life.

5 Conclusion: Back to Computational Creativity

In presenting this review of computational Artificial Life, and particularly the evolution of complexity, I have suggested that this can be usefully related to more general problems of computational creativity. In concluding then, it is appropriate to return to this claim, and at least outline a basis for such connections.

I take the view that the form of complexity discussed above is essentially equivalent to genetically recorded *knowledge*: that is, systems grow in complexity to the extent that they develop internal predictive models of their worlds, which allow them to anticipate, plan, and interact, in evolutionarily effective ways. In this sense, Darwinian evolution is a process of knowledge growth, typically taking place over geological timescales and using DNA as an information repository. The knowledge which results is realised in the genetically constrained development, morphology, and behaviours of organisms; and it is, manifestly, highly context-dependent and fallible.

A natural question is how this growth of genetically recorded knowledge relates to the more familiar, somatic time, growth of knowledge in an individual human mind, and its expression, communication and collective growth in a human culture; and this is relevant because the latter is surely a characteristic example of human creativity.

The nature of our knowledge of the world, or epistemology in general, is, of course, a foundational problem in philosophy. I will not attempt to reprise millennia of debate here but will simply advocate one particular position. This is the *evolutionary epistemology* pioneered by Popper (1976) and Campbell (1974). This holds that our individual, learned, human knowledge, and our accumulated human cultural knowledge is fundamentally no different from our biologically derived evolutionary, genetic

knowledge. That is, human knowledge is also irredeemably context-dependent, and ultimately fallible; and insofar as it grows, both for an individual and an entire human civilization, it does so by a process formally identical with that realised by biological evolution, namely unjustified variation and selective retention.

While this is by no means a consensus position in modern epistemology, it does enjoy a reasonably broad level of support. In particular, it is the foundation for the *Universal Selection Theory* of Cziko (1995); and provides a philosophical basis for Edelman (1987)'s *Neural Darwinism* and Calvin (1996)'s *Neural Darwin Machine*, both of which address the realisation of essentially Darwinian evolutionary processes in the somatic time development and behaviour of complex nervous systems.

So if all of this can be accepted, then computational realisations of Darwinian processes, which demonstrate spontaneous and sustained growth of complexity, would constitute the strongest examples which I can currently imagine of autonomous computational creativity; and in doing so, would surely shed much light on deep aspects of human creative activities. Of course, while I have summarised much interesting work on these problems, the outstanding difficulties, both practical and theoretical, remain deep seated. Nonetheless, I hope that I have demonstrated the substantial confluence of interest between computational evolution and computational creativity, and the potential for continuing fruitful interactions between researchers from these two communities.

References

- Adami, C. and Brown, C. T. (1994). Evolutionary learning in the 2D artificial life system "Avida". In Brooks, R. A. and Maes, P., editors, *Proc. Artificial Life IV*, pages 377–381. MIT Press. Available from: <http://citeseer.ist.psu.edu/101182.html>.
- Arthur, S. (1959). Some studies in machine learning using the game of checkers. *IBM Journal*, 3(3):210–229. Available from: <http://www.research.ibm.com/journal/rd/033/ibmrd0303B.pdf>.
- Axelrod, R. (1984). *The Evolution of Cooperation*. Basic Books.
- Axelrod, R. (1987). The evolution of strategies in the iterated prisoner's dilemma. In Davis, L., editor, *Genetic Algorithms and Simulating Annealing*. Pitman.
- Banzhaf, W., Nordin, P., Keller, R., and Francone, F. (1998). *Genetic Programming - An Introduction*. Morgan Kaufman.
- Barricelli, N. (1957). Symbiogenetic evolution processes realized by artificial methods. *Methodos*, 9(35-36):143–182.
- Barricelli, N. (1963). Numerical testing of evolution theories. *Acta Biotheoretica*, 16(3):99–126.

- Berlekamp, E. R., Conway, J. H., and Guy, R. K. (1982). What is life? In *Winning Ways for your Mathematical Plays*, volume 2, chapter 25, pages 817–850. Academic Press, London.
- Bratley, P. and Millo, J. (1972). Computer recreations; self-reproducing automata. *Software—Practice and Experience*, 2:397–400.
- Brown, P. (1996). Emergent behaviours: towards computational aesthetics. *Artlink*, 16(2 & 3). Available from: <http://www.paul-brown.com/WORDS/EMERGENT.HTM>.
- Burks, A. W., editor (1966). *Theory of Self-Reproducing Automata [by] John von Neumann*. University of Illinois Press, Urbana.
- Calvin, W. H. (1996). *The Cerebral Code: Thinking a Thought in the Mosaics of the Mind*. MIT Press. Available from: <http://www.williamcalvin.org/bk9/>.
- Campbell, D. T. (1974). Evolutionary epistemology. In Schilpp, P. A., editor, *The Philosophy of Karl Popper*, volume XIV of *The Library of Living Philosophers*, pages 413–463 (Book I). Open Court, Illinois.
- Cariani, P. (1992). Some epistemological implications of devices which construct their own sensors and effectors. In *Toward a practice of autonomous systems-Proceedings of the First European Conference on Artificial Life*, pages 484–493.
- Chu, D. and Ho, W. K. (2006). A category theoretical argument against the possibility of artificial life: Robert Rosen’s central proof revisited. *Artificial Life*, 12(1):117–134. Available from: <http://www.mitpressjournals.org/doi/abs/10.1162/106454606775186392>.
- Codd, E. F. (1968). *Cellular Automata*. ACM Monograph Series. Academic Press, Inc., New York.
- Cziko, G. (1995). *Without Miracles: Universal Selection Theory and the Second Darwinian Revolution*. MIT Press. Available from: <http://faculty.ed.uiuc.edu/g-cziko/wm/>.
- Dewdney, A. K. (1984). Computer recreations: In a game called Core War hostile programs engage in a battle of bits. *Scientific American*, 250:14–22.
- Dewdney, A. K. (1987). Computer recreations: A program called mice nibbles its way to victory at the first core wars tournament. *Scientific American*, 256(1):8–11.
- Dewdney, A. K. (1989). Computer recreations: Of worms, viruses and core war. *Scientific American*, 260(3):90–93.
- Dittrich, P. and di Fenizio, P. S. (2007). Chemical organization theory. *Bull. Math. Biol.*, 69(4):1199–1231.

- Dittrich, P., Ziegler, J., and Banzhaf, W. (2001). Artificial Chemistries - A Review. *Artificial Life*, 7(3):225–275.
- Dobzhansky, T. (1973). Nothing in biology makes sense except in the light of evolution. *The American Biology Teacher*, 35:125–129. Available from: http://www.pbs.org/wgbh/evolution/library/10/2/text_pop/1_102_01.html.
- Dyson, G. (1997). Darwin among the machines; or, the origins of [artificial] life. Presentation hosted by the Edge Foundation, Inc., 8 July 1997. Available from: http://edge.org/3rd_culture/dyson/dyson_p1.html.
- Dyson, G. B. (1998). *Darwin Among the Machines: The Evolution of Global Intelligence*. Basic Books.
- Edelman, G. M. (1987). *Neural darwinism*. Basic Books, New York.
- Farmer, J., Kauffman, S., and Packard, N. (1986). Autocatalytic replication of polymers. *Physica D*, 22:50 – 67.
- Farmer, J. D. and d’A. Belin, A. (1992). Artificial life: The coming evolution. In Langton et al. (1992), pages 815–838.
- Fontana, W. and Buss, L. (1994). The arrival of the fittest: Toward a theory of biological organization. *Bulletin of Mathematical Biology*, 56:1 – 64.
- Hofstadter, D. R. (1979). *Gödel, Escher, Bach: An Eternal Golden Braid*. Penguin Books, London. First published in Great Britain by The Harvester Press Ltd 1979. Published in Penguin Books 1980.
- 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. Proceedings of a conference held in Noordwijkerhout, Holland, 31 March – 6 April 1975.
- Kampis, G. (1991). *Self-Modifying Systems in Biology and Cognitive Science*, volume 6 of *IFSR International Series on Systems Science and Engineering*. Pergamon Press, Oxford. Editor-in-Chief: George J. Klir.
- Kauffman, S. A. (1993). *The origins of order: self-organization and selection in evolution*. Oxford University Press, Oxford.
- Koza, J. (1992). *Genetic Programming*. MIT Press.
- Kvasnicka, V., Pospíchal, J., and Kaláb, T. (2001). A study of replicators and hypercycles by typogenetics. In *ECAL 2001: Proceedings of the 6th European Conference on Advances in Artificial Life*, pages 37–54. Springer-Verlag.
- Langton, C. G. (1984). Self-reproduction in cellular automata. *Physica*, 10D:135–144.

- Langton, C. G., editor (1989). *Artificial Life*, volume VI of *Series: Sante Fe Institute Studies in the Sciences of Complexity*. Addison-Wesley Publishing Company, Inc., Redwood City, California. Proceedings of an interdisciplinary workshop on the synthesis and simulation of living systems held September, 1987, in Los Alamos, New Mexico.
- Langton, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S., editors (1992). *Artificial Life II*, volume X of *Series: Sante Fe Institute Studies in the Sciences of Complexity*. Addison-Wesley Publishing Company, Inc., Redwood City, California. Proceedings of the workshop on Artificial Life held February, 1990, in Sante Fe, New Mexico.
- McMullin, B. (1992). The Holland α -universes revisited. In Varela, F. J. and Bourgine, P., editors, *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, Series: Complex Adaptive Systems, pages 317–326, Cambridge. MIT Press. Series Advisors: John H. Holland, Christopher Langton and Stewart W. Wilson.
- McMullin, B. (2000a). John von Neumann and the evolutionary growth of complexity: Looking backwards, looking forwards. *Artificial Life*, 6(4):347–361. Available from: <http://www.eeng.dcu.ie/~{}alife/bmcm-alj-2000/>.
- McMullin, B. (2000b). Some remarks on autocatalysis and autopoiesis. *Annals of the New York Academy of Sciences*, 901:163–174. Available from: <http://www.eeng.dcu.ie/~{}alife/bmcm9901/>.
- McMullin, B. (2004). 30 years of computational autopoiesis: A review. *Artificial Life*, 10(3):277–296. Available from: <http://www.eeng.dcu.ie/~{}alife/bmcm-alj-2004/>.
- McMullin, B., Taylor, T., and von Kamp, A. (2001). Who needs genomes? Presented at the Atlantic Symposium on Computational Biology, Genome Information Systems & Technology March 15-17, 2001, Regal University Center, Durham, N.C., USA. Available from: <http://www.eeng.dcu.ie/~alife/bmcm-cbgi-2001/>.
- McMullin, B. and Varela, F. J. (1997). Rediscovering computational autopoiesis. In Husbands, P. and Harvey, I., editors, *Proceedings of the Fourth European Conference on Artificial Life (ECAL-97)*, Series: Complex Adaptive Systems, Cambridge. MIT Press. Available from: <http://www.eeng.dcu.ie/~{}alife/bmcm-ecal97/>.
- Pargellis, A. N. (2001). Digital life behavior in the amoeba world. *Artificial Life*, 7(1):63–75. Available from: <http://www.mitpressjournals.org/doi/abs/10.1162/106454601300328025>.
- 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.

- Rasmussen, S., Knudsen, C., Feldberg, R., and Hindsholm, M. (1990). The coreworld: Emergence and evolution of cooperative structures in a computational chemistry. *Physica*, 42D:111–134.
- Ray, T. S. (1992). An approach to the synthesis of life. In Langton et al. (1992), pages 371–408.
- Ray, T. S. (1995). A proposal to create two biodiversity reserves: One digital and one organic. Project proposal for Network Tierra. Available from: <http://life.ou.edu/pubs/reserves/>.
- Rocha, L. M., editor (2000). *The Physics and Evolution of Symbols and Codes: Reflections on the Work of Howard Pattee*. Elsevier. Special Issue of BioSystems, Vol. 60, No. 1-3.
- Rosen, R. (1959). On a logical paradox implicit in the notion of a self-reproducing automaton. *Bulletin of Mathematical Biophysics*, 21:387–394.
- Rosen, R. (1972). Some relational cell models: The metabolism-repair systems. In Rosen, R., editor, *Foundations of Mathematical Biology*, volume II. Academic Press, New York.
- Rosen, R. (1985). Organisms as causal systems which are not mechanisms: An essay into the nature of complexity. In Rosen, R., editor, *Theoretical Biology and Complexity*, chapter 3, pages 165–203. Academic Press, Inc., Orlando.
- Rosen, R. (1991). *Life Itself*. Columbia University Press, New York.
- Rozenberg, G., Bäck, T. H., and Kok, J. N., editors (2010). *Handbook of Natural Computing*. Springer. Available from: <http://www.springer.com/computer/foundations/book/978-3-540-92911-6>.
- Satomi, M. and Sommerer, C. (2007). "game_of_life": interactive art installation using eye-tracking interface. In *ACE '07: Proceedings of the international conference on Advances in computer entertainment technology*, pages 246–247, New York, NY, USA. ACM.
- Snare, A. (1999). Typogenetics. Thesis for Bachelor of Science (Computer Science) Honours, School Of Computer Science And Software Engineering Monash University. Available from: <http://www.csse.monash.edu.au/hons/projects/1999/Andrew.Snare/thesis.pdf>.
- Thatcher, J. W. (1970). Universality in the von neumann cellular model. In Burks, A. W., editor, *Essays on Cellular Automata*, pages 132–186 (Essay Five). University of Illinois Press, Urbana.
- Turing, A. (1936). On computable numbers, with an application to the entscheidungsproblem. *Proceedings of the London Mathematical Society*, Series 2, Vol. 42:230–265.

- 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. (A corrected version was formally published in the IEEE Annals of the History of Computing, 15(4), 1993.).
- von Neumann, J. (1949). Theory and organization of complicated automata. In Burks (1966), pages 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.
- Vyssotsky, V. A. (1972). Darwin: A game of survival and (hopefully) evolution. *Software—Practice and Experience*, 2:91–96. Attachment to letter appearing in Computer Recreations column, signed by M. D. Mcilroy, R. Morris and V. A. Vyssotsky. Available from: <http://www.cs.dartmouth.edu/~doug/darwin.pdf>.
- Willis, M., Hiden, H., Marenbach, P., McKay, B., and Montague, G. A. (1997). Genetic programming: An introduction and survey of applications. In Zalzala, A., editor, *Second International Conference on Genetic Algorithms in Engineering Systems: Innovations and Applications, GALESIA*, pages 314–319, Savoy Place, London WC2R 0BL, UK. Institution of Electrical Engineers (IEE). Available from: <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=00681044>.