School of Computer Science & Software Engineering Bachelor of Software Engineering Honours Literature Review, Semester 2, 2003 Evolving Non-Standard Models of Computation Supervisor: A/Prof David Dowe

Stephen Jackel,
Student ID: 12782548,
Monash University, Clayton, Australia
Email: stjac1@student.monash.edu.au
Computer Science & Software Engineering
http://www-personal.monash.edu.au/ stjac1/honours.html

July 30, 2003

Contents

1 Introduction

1.1 Is there life on other planets?

At present we are aware only of life on planet Earth. Many people argue that Earth-like planets may be extremely rare. Thus, assuming that life requires Earth-like planets, life on Earth may be unique. Others propose that in a large enough universe, possibly with many Earth-like planets, that life elsewhere is inevitable [30, 17]. But, the form of this life is unknown. Indeed, no-one is sure that other life will even display signs of intelligence. In fact, there is even debate into the exact nature of life and intelligence. This makes recognising and measuring intelligence in other beings very difficult. Despite efforts to find life on other planets, no conclusive evidence has been found, and given the massive distances to even our closest neighbour stars, we may never be able to solve this debate with certainty [3]. Thus the fields of artifical intelligence and artificial life attempt to find answers by turning to life in the digital medium. Intelligent organisms have been evolved in artificial life simulators. However, the best methods of encouraging organisms to learn and the computations that they will be capable of performing remain uncertain. This report examines research into these areas and past attempts to understand where we have come from and what intelligence may exist in our universe.

2 What is Life?

2.1 Definition of life

What is life? Still no-one is able to give a complete answer to this age old question. The question arose again in a slightly different form in the early 1950's as the rise of digital computing prompted the question, "What is intelligence?" Could computers be capable of displaying intelligence? Could we create "artificial" intelligence? Would that intelligence be considered alive? In 1950, the Imitation Game, later known as the Turing Test was suggested by Alan Turing as a way of defining intelligence [35]. Intelligent behaviour was defined as being the ability to act in a (restricted) manner such that an observer could not tell if the entity was human or not. Since then, the Turing Test has been used as a guide to both life and intelligence in the field of computer science [23]. But life and intelligence, while perhaps related, are very different. It is easy to imagine or find living organisms that exhibit intelligence but would fail the Turing Test. Even humans who speak another language may have great difficulty. I, and many others, believe that the Turing Test is more useful as a benchmark and definition only of human intelligence rather than of the intelligence of any entity [6, 10, 13]. Therefore, we must find other means of defining and testing for life and intelligence when searching in the digital realm.

2.2 Artificial Life

Life forms as we know them are based in the medium of carbon chemistry. Assuming that the Darwinian theory that life on Earth evolved through the process of natural selection is correct, we imagine that life on other planets would be similar. But we have only seen an instance of life on Earth; life may take many forms. Biology and an attempt to understand life on our planet helped to spawn the field of artificial life, where computer programs attempt to simulate evolution and/or grow artificial life in the digital medium with the general goal of achieving life of comparable complexity to our own [29, 21].

However, there has been significant debate as to whether life produced in other media is just a simulation of life in the physical world or whether this other life is a true realisation of life in its own right. Some argue that the physical medium is crucial to life and that true life is not possible in other media [11, 22]. Others, including Thomas Ray, believe that life is a process or set of processes that may be implemented in any physical medium that is able to facilitate these processes [16, 25]. Although this issue is far from resolved, I tend to agree with those who believe that some sort of physical environment with its own "physics" is required for life and that it is likely that the symbolic environment of a computer program in the digital medium is able to provide this [24, 5]. Indeed, it has been suggested that the carbon-based medium in which we live is itself a simulation [4].

Regardless of whether or not digital life is "real" life, the digital medium as implemented on computers is an attractive place to simulate life as countless variables and values of experiments can be measured and analysed, and experiments can be repeated exactly, unlike in biology where slight variations in conditions or hidden variables may cause different results. Vast improvements in computational power and storage space have enabled researchers to watch evolution in "fast-forward".

While many simulations have been run to gain insights into evolution in our world, the results are only useful as far as the model that implements them will stretch; the more detailed and realistic the model, the more accurate the results. But forcing the digital medium to imitate our carbon-based world of chemistry and physics becomes increasingly difficult as more detail is added [29]. Although much insight may be gained through such experimentation, organisms that evolve on such systems will never be the same as in the real world because they are not in the real world; they are influenced by the medium in which they live.

The digital medium is almost completely independent of the carbon-based one in which we live; it has its own rules, its own chemistry and physics. Organisms that evolve in the digital medium will adapt to and belong in the digital medium. This new environment is likely to produce very different life and intelligence to that with which we are familiar. Thomas Ray believes that we will never be able to guide evolution into producing a particular form due to lack of foresight - the "algae to corn" problem [28]. For example, having never seen anything more complex than single-celled algae, it would be virtually impossible to have the insight to devise a fitness function that would guide the organisms through many levels of complexity to produce corn. Transferring this across to the digital medium, we have no idea where digital intelligence may go or what forms it will take and thus we should not try to impose too many fitness functions or other restrictions, but rather let evolution find its own path. While I believe that this is true I suspect that it will be possible to teach organisms skills that *may* be useful in guiding their paths.

3 Using Digital Evolution to Solve Problems

Biologists have been attempting to recreate biological life in digital form for a number of years. Nevertheless, there have been other attempts to harness the power of evolution. Genetic algorithms (GA) have been used to find optimal solutions to problems [20]. GAs are limited, however, as the form of the solution must already be known; GAs merely attempt to find the best values of the solution. This method is thus not useful in discovering new forms of intelligence or computation. An improvement comes in the form of Genetic Programming (GP), where the form of the solution is subject to change according to a pre-defined fitness function. Genetic programming appears to be more likely to find new forms of intelligence but, as predicted by Thomas Ray, it is limited to the fitness function. A fitness function may provide a strong guiding force for evolution to begin with. However, unless it is extremely well thought-out, it will become weaker as the organisms grow towards it until the fitness function becomes a ceiling beyond which evolution cannot pass.

Fitness functions that change as organisms evolve have been used to counter this problem, the most successful being co-evolution [20]. Co-evolution involves the fitness function also being subject to evolutionary processes rather than being fixed or embedded in the medium. The organisms grow towards satisfying the fitness function while the fitness function tries to maximise the gap between itself and the level of the evolving organisms. Effectively, instead of the fitness function being a fixed bar, the bar is raised as the organisms evolve closer. However, this is only a successful method if the fitness function stays ahead (but not by too much) of the evolution of the organisms. Finally, at the far extreme is Thomas Ray's idea of letting the medium dictate the selective forces, and observing what happens. Thomas Ray has attempted to implement this idea in Tierra [26, 28].

3.1 Tierra

Tierra is an Artificial Life (AL) simulator written by Thomas Ray that attempts to use evolution in the "digital medium" to bring forth life (usually inoculated with a simple reproducing program) [26]. Written in C, Tierra implements a virtual computer complete with "Darwinian" operating system and architecture that allows machine code to be evolved through randomly changing bits, "mutation", and interchanging sections of code, "recombination". The system supports evolving programs or "organisms" and overcomes the "brittleness problem" - the fact that most random changes to valid code will produce invalid code. Tierra allows for the inevitable lines and sections of code that are not executable by ignoring them. Over time, beneficial changes are made to the code that enable the digital organisms living in the system to evolve and increase in function and complexity.

There are a number of control and monitoring tools built into the simulator. Tierra records all births and deaths and keeps a "genebank" of all successful genomes (sequences of creature code) and is able to monitor and record creature interactions. In principle, Tierra provides an environment that is able to foster parallel learning. This means that the time taken for organisms to learn to perform several tasks should be less that the time required to learn each of the tasks sequentially. Thus Tierra is an extremely useful tool with which to observe, measure and experiment with evolution in the digital medium.

Organisms in the digital world of Tierra face the obstacles of various inherent natural limitations and finite resources. The RAM available on the computer running the simulation forms the space or "soup" that the organisms must occupy and share. Likewise, in the digital medium, the CPU is likened to energy as it is the only "force" that allows creatures to change state, reproduce and perform other actions and is therefore in high demand [28]. Competition for these resources creates evolutionary pressures for the organisms to adapt to their environment and outperform their competitors.

The instructional set used by the CPU forms the “physics and chemistry” of the digital medium together with the methods used to allocate resources. In Tierra, the instructional set is deliberately designed to occasionally introduce errors in calculations, feeding evolution. Organisms in memory may read or execute instructions from anywhere (or in some implementations, within a certain proximity of an organism) but may only write to memory that it “owns”. These rules help to define the ways in which digital organisms may interact with others.

Organisms age through the use of a “reaper queue” in which the organisms at the front of the queue (usually the oldest organisms) are killed to make way for new organisms [25]. However, organisms that attempt to perform illicit operations, such as writing to the memory space occupied by another organism, are accelerated through the queue. This is done by regularly comparing the number of illegal operations performed by an organism with the number performed by the organism in front of it in the reaper queue. If the organism has performed more illicit operations than the other organism, the two switch places in the queue. This is a way of punishing organisms that are not abiding by the “rules” of the world and constitutes a “natural” fitness function provided by the Tierra environment. It is also possible to reward organisms that are performing “good” operations by moving such organisms backward one place in the reaper queue, meaning that the organism will live for longer. This is one way in which the user may implement a fitness function that guides organisms by encouraging them to perform certain operations. Distributing bonus CPU time to organisms that perform certain actions is another method of providing a fitness function in Tierra.

Thomas Ray has produced an array of different organisms in Tierra by inoculating the simulation with a simple creature capable only of reproducing itself [26]. Thus all other knowledge or skills that the organisms developed had to evolve (through mutations or other means). Eventually, mutations were not required as the creatures learnt to sexually reproduce, having offspring that were a mixture of the parents. Thomas Ray, pointing to the Amazon rain forests as another example, argues that as organisms increase in complexity, the interactions between the organisms often form a much stronger evolutionary force than the medium upon which they are implemented [28].

3.2 Extensions to Tierra

Tierra was extended to “Network Tierra” through the use of a global “Digital Reserve” [27]. The idea was that organisms would be able to roam over a form of the Internet, migrating between different computers on the network in search of resources. People were able to donate CPU time in a similar manner to SETI [1]. It was hoped that this expansion in the digital universe and introduction of further, complex interactions would set about an explosion of multi-cellular digital life equivalent to the explosion of multi-cellular life on earth thought to have occurred during the Cambrian period [36]. However, the amount of traffic generated over the networks of participating organisations was larger than expected and placed a substantial strain on both the finances and the network bandwidth of the participating organisations. A cluster model with reduced bandwidth was tried but failed to produce any interesting results [29].

Tierra is among the most widely known Artificial Life simulators, though there are a number of others, such as Avida. Avida is a Tierra derivative that allows the introduction of fitness functions in addition to the natural ones intrinsic in the medium and implementation [21]. All artificial life simulations may be viewed as tools that allow us to gain an insight into our own world and the life around us and make educated predictions about life elsewhere in the universe.

4 Search For Non-Standard Models of Computation

Hand-held calculators are able to perform mathematical computations. Humans, too, are able to perform many different computations, some of which are the same as a hand-held calculator. Hand-held calculators that are working correctly will always perform a certain calculation in the same amount of time and will always produce the same answer. However, the method that humans use, especially when done without the use of paper and pens or other tools, is somewhat haphazard and the same calculation may even produce different results on different occasions. Likewise, calculators must deal with finite values whereas humans can perform calculations that include abstract concepts such as infinity. Clearly different models of computation are being used. While some methods of performing computations are much faster and more efficient than others, some are also more powerful and are able to compute a larger range of problems. Models can be tested for Turing Completeness; if a model is Turing Complete then it is able to compute any computable function. It is possible that other models of computation are much more efficient or even more powerful than those currently in use [6]. For example, Jack Copeland has suggested an “accelerated Turing Machine” that is theoretically able to compute uncomputable functions, such as the halting function [7]. Research into quantum computing and biomolecular computing is also attempting to find new models of computation [8, 19, 20, 32].

While many of these areas hold great promise, others turn to the intelligence and capability to compute seen in living entities. Humans are intelligence beings (at least compared with all creatures that we have discovered). Humans have used their intelligence to create machines: cars, hand-held calculators, printing presses and computers. None of these machines are considered intelligent themselves, rather they have had intelligence imbedded in them by humans. But what about other entities, whether aliens, dolphins or digital agents? Do/would they display signs of intelligence? Their intelligence and abilities may be very different to our own [34, 10], so how would we measure it? Discovering and measuring their intelligence or computational ability is an issue that many have considered, but still there are few solid answers.

4.1 What is intelligence?

Intelligence is not an easy thing to define. It is commonly agreed that systems are intelligent if they use all available and as much previously available knowledge as necessary when solving a problem [37]. Lack of knowledge is different to a lack of intelligence. Therefore it is best to compare two systems which both have access to the same knowledge - i.e. imbed the information required to solve a puzzle in the problem. But this is more of an observation than a definition. The Merriam-Webster Dictionary defines intelligence as "(1) the ability to learn or understand or to deal with new or trying situations (2) the ability to apply knowledge to manipulate one's environment or to think abstractly as measured by objective criteria (as tests)". John McCarthy, commonly known as the father of artificial intelligence, states that "Intelligence is the computational part of the ability to achieve goals in the world" and adds that intelligence can be seen in humans, many animals and even some machines [18]. While these are reasonable descriptions of intelligence, there is no full answer; it is an elusive concept.

4.2 Collective Intelligence

The idea that a group of entities might display an overall intelligence greater than any of the individual entities is known as collective intelligence and is related to both physical synergy and problem-solving that involves a number of individuals. It is an area that encompasses such organisations as ant colonies, human populations, bacterial colonies and even groups of molecules. Hofstadter considered the relationship between a human mind and an ant colony in that while each individual part (ant or neuron) may not have much intelligence, collectively the intelligence of the mind or ant colony is relatively large [15]. Szuba has used the Random PROLOG Processor to model collective intelligence in social structures such as bacterial and insect colonies and the model extends to include human social structures [33]. Together with Almulla, Szuba went on to question the links between intelligence and life [34]. They reasoned that collective intelligence can be found in both life and in non-living groups of molecules. As the requirements for collective intelligence are weaker than for life, perhaps collective intelligence appeared on Earth in collections of molecules before life emerged. Unfortunately, no definite conclusion is reached as to exactly how collective intelligence emerges and Szuba and Almulla's proposed relationship between life and intelligence cannot be considered any more than a draft hypothesis at this stage. Thus there is still much to learn about this powerful and interesting form of intelligence and method by which computations may be done.

4.3 Measuring Intelligence and Ability

The question of how to measure intelligence and an entity's computational or problem solving abilities is also an area of hot debate, running to the core of the field of artificial intelligence. Intelligence Quota (IQ) tests are used by humans and appear to be the most obvious choice. However, many believe that IQ tests are not suitable and/or flexible enough to be used to measure the intelligence of entities other than humans [10, 12, 14, 31].

Most accept that it is possible to determine the intelligence of an agent if the agent's behaviour and knowledge are known [37]. As Hajek and Dowe explain, the agent should be tested as a whole but the internal workings or architecture of the agent must also be analysed [9, 31]. Gudwin agrees, and believes that both black and white box testing must be considered in order to evaluate an agent's intelligence [14]. Gudwin suggests that the type of architecture should be placed into one of six categories. Then, by classifying types of knowledge into a hierarchy, intelligence can be given by a pair of coordinates into a lattice (X,Y), where X = type of architecture, and Y = the number of different types of knowledge the system can handle. However, I don't believe that this system compares the levels of knowledge that different agents may have fairly. It is also based upon the assumption that the internal workings of the agent are both available and are understood. It is often difficult to specify all of the agent's knowledge which then must be inferred, perhaps from tests. To avoid this problem, Behavioural Bounding together with a general algorithm can be used to determine an agent's intelligence [37]. However, it may require impractical amounts of computational resources to run and in practice this technique may be difficult to apply.

Perhaps a more general type of IQ test based on the general intelligence factor and the theory of multiple types of intelligence is the answer [12]. These could include testing an agent's performance and response in various known and unknown environments and situations. However, while useful guidelines are available for the criterion of ideal intelligence tests, it is not always evident how these may be implemented and how it can be known if a complete or fair set is being used. It appears that the exact nature of intelligence and the manner in which intelligence can be measured and fairly compared between different types of agents is as elusive as ever.

4.4 Evolving Intelligence and the Ability to Perform Computations in Tierra

Chris Adami used Tierra to investigate the evolution of intelligence and the ability of self-replicating digital organisms to learn to perform computations. He ran a number of experiments to examine the learning capabilities of the organisms in the hope of discovering some universal features of the learning process [2]. The general learning characteristics of organisms in Tierra are thought to be independent of the particular learning task, making the choice of learning task somewhat arbitrary. Adami also experimented with how fast organisms could reach a certain level of intelligence by altering the mutation rate of the simulation.

A number of simulations were run in which evolving organisms learnt to compute the sum of two numbers. This was done through the use of a three-stage fitness function which rewarded or punished organisms by varying the allocated CPU time to each organism whenever they were successful or unsuccessful in performing certain tasks.

The number of offspring that an organism produced in a certain time period (where the time period was a certain number of instructions executed) was found to be a useful measure of the fitness of the organism. As organisms were rewarded for performing "intelligent" operations, the fitness of an organism was related to it's intelligence. By averaging this value over all organisms in the simulator, a universal measure of the overall fitness or intelligence of the system was estimated. Often there were periods where the fitness of the system remained fairly constant as organisms were in a sort of equilibrium. After some time there was often a "break-through" as some organisms evolved a new ability that increased their fitness. This new ability or increase in intelligence usually spread through the system and the overall fitness of the system rose noticeably before reaching another plateau.

After running a set of simulations, the mutation rate was modified and another set of simulations were run. The results showed that the amount of time required for a simulation to reach the goal varied greatly even when the starting conditions (with the exception of the random seed) were identical. It was also found that increasing the mutation rate increased the rate at which the overall fitness of the system improved and thus increased the learning rate of the organisms. However, if the mutation rate went above a certain threshold, organisms were not able to reproduce before they were subjected to further mutation. This situation rapidly led to a "melt-down" and extinction of organisms in the system. Thus the environmental variables do help to determine the rate of intelligence growth, which is highest when the random mutations are almost too frequent for life to be maintained. Adami also defined a "learning fraction" that effectively gives the probability of a particular Tierra simulation run producing organisms that have the ability to add within a certain time limit (amount of CPU time). It is possible that future research will uncover a learning fraction can be generalised to predict the probability of organisms reaching a certain level of intelligence within a certain amount of CPU time.

5 Conclusion

In conclusion, the concept of life has been examined and has been found to exist on Earth and quite possibly elsewhere in the universe. Artificial life in the digital medium may provide us with a way of exploring life and intelligence. Digital evolution has indeed already been used to solve a number of real-world problems and provide insight into how our carbon-based universe may operate. The Tierra simulator has pioneered much research into these areas and has been used to successfully evolve a variety of digital organisms. After considering the search for non-standard models of computation, the possibility of evolving an intelligence capable of performing computations was examined. Different forms of intelligence such as collective intelligence were studied along with the inherent difficulties in both defining and measuring intelligence. Finally, research by Chris Adami was detailed in which he experimented with evolving and measuring the intelligence of Tierra organisms and the factors affecting the rate at which the organisms learnt to perform basic computations.

5.1 Relevance of my Research

My project looks at evolving and "teaching" digital organisms. It is hoped that insights gained from the field of artificial life will help us understand life on our planet and perhaps be a guide to the possibility of life elsewhere in our universe.

I will modify Tierra to allow a specialised sort of fitness function (or "teaching function") to regulate the allocation of CPU time to the digital organisms. The goal will be to try to nurture intelligence so that the organisms are able to perform computations, quite possibly in a manner that we have not seen before. Thus the project attempts to find a balance between allowing the digital organisms enough freedom in their natural digital medium to evolve and display complex behaviour, and causing the organisms to display behaviour that is useful to us in the real world. That is, we want to be able to see and measure what computations the digital organisms are able to perform. I would like to be able to gain some insight into their "intelligence" by seeing how they react to new calculations that are similar to those that they can already solve. I hope to evolve digital organisms that are capable at least of the most basic mathematical computations.

My project will be beneficial in the quest to understand life on Earth and how intelligence might evolve. It will contribute to the search for non-standard models of computation by attempting to discover the potential of finding new models of computation through evolution in a digital world. It will also examine how practical it is to guide evolution into performing computations and whether or not teaching functions will speed-up the evolutionary process. A teaching function may be a method of providing a ratchet mechanism to prevent degeneration of learnt ability, thus allowing the creation of building blocks that allow organisms to move up the evolutionary chain towards a Digital Cambrian explosion of complexity. Will the digital organisms be able to use their new knowledge for new purposes? Is this a way of evolving intelligence? Are such teaching methods are a useful way of guiding evolving organisms towards an end goal or at least speeding up the evolutionary process? Thus my project will build upon the existing knowledge that has been considered in this report and will probe into new methods of evolving non-standard models of computation.

References

[1] The Search for Extraterrestrial Intelligence (SETI). http://setiathome.ssl.berkeley.edu/ There is a large amount of information about the SETI project available on the website, including free download of the SETI at Home software. David Anderson was the founder of SETI at Home.

[2] Chris Adami. Learning and Complexity in Genetic Auto-Adaptive Systems. *Physica D*, 80:154–170, 1995. http://xxx.lanl.gov/PS_cache/adap-org/pdf/9401/9401002.pdf OR for a postscript version visit http://xxx.lanl.gov/abs/adap-org/9401001.

[3] J. R. P. Angel and N. J. Woolf. Searching for Life on Other Planets. *Scientific American*, 274(4):60–66, 1996. http://atropos.as.arizona.edu/aiz/teaching/a204/etlife/SciAm96.pdf.

[4] John D. Barrow. Glitch! How our World could be just a Computer Simulation. *New Scientist*, 7 June, 2003.

[5] Mark A. Bedau. Four Puzzles About Life. *Artificial Life*, 4(2):125–140, 1998.

[6] C. S. Calude and M. J. Dinneen. Breaking the Turing Barrier. Technical report, 1998. CDMTCS Research Report 084, 5pp. http://citeseer.nj.nec.com/calude98breaking.html.

[7] Jack B. Copeland. Even Turing Machines can Compute Uncomputable Functions. *Unconventional Models of Computation*, 1998. Springer, Editors: C.S. Calude, J. Casti and M.J. Dinneen, Paper presented at the First International Conference on Unconventional Models of Computation, University of Auckland (City Campus), Auckland, New Zealand, UMC.

[8] Rajarshi Das. Emergent Computation in Cellular Automata, 1998. Centre for Nonlinear Studies, MS-B258, Los Alamos National Laboratory, Los Alamos. http://citeseer.nj.nec.com/46929.html.

[9] D.L. Dowe and A.R. Hajek. A Non-behavioural, Computational Extension to the Turing Test. 1998. Proceedings of the International Conference on Computational Intelligence and Multimedia Applications (ICCIMA9́8), Gippsland, Australia, February, pp101-106.

[10] Bruce Edmonds. The Social Embedding of Intelligence: Towards Producing a Machine that could Pass the Turing Test, 2002. To be published as Edmonds, B. (in press) The Social Embedding of Intelligence - Towards a Machine that could Pass the Turing Test. In Peters, G. and Epstein, R (eds.) The Turing Test Sourcebook: Philosophical and Methodological Issues in the Quest for the Thinking Computer, Kluwer. http://citeseer.nj.nec.com/534517.html OR http://cfpm.org/cpmrep95.html.

[11] Claus Emmeche. Life as an Abstract Phenomenon: Is Artificial Life Possible? 1992. Towards a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life, pages 466-474.

[12] Dino Franklin and Ana Abrao. Measuring Software Agent's Intelligence. 2000. SSGRR 2000, International Conference on Advances in Infrastructure for Electronic Business, Science and Education on the Internet, LÁquila, Brasil, July-August, 2000. Authors from Federal University of Uberlandia, Brasil. http://citeseer.nj.nec.com/452553.html OR http://www.ssgrr.it/en/ssgrr2000/papers/209.pdf.

[13] R. M. French. Subcognition and the Limits of the Turing Test. *Mind*, 99:53–65, 1990.

[14] Ricardo R. Gudwin. Evaluating Intelligence: A Computational Semiotics Perspective. 2000. IEEE International Conference on Systems, Man and Cybernetics, SMC2000, Nashville, Tenessee, USA, October, 2000, pages 2080-2085. http://www.dca.fee.unicamp.br/ gudwin/ftp/publications/smc2000.pdf.

[15] Douglas Hofstadter. Godel Escher Bach: An Eternal Golden Braid. 1979. New York: Basic Books, pp. 275-336. Particularly note the section flanking chapter 10.

[16] Christopher G. Langton. Studying Artificial Life with Cellular Automata. *Physica D*, 22:120–149, 1986.

[17] Betsy Mason. Look of Life... *New Scientist*, 12 July:28–31, 2003.

[18] John McCarthy. *What is Artificial Intelligence?* PhD thesis, Stanford University, 2000. Originally published April 4, 2000 as an academic paper. http://www.kurzweilai.net/articles/art0088.html?printable=1 OR http://www-formal.stanford.edu/jmc/whatisai.html.

[19] Melanie Mitchell, James P. Crutchfield, and Peter T. Hraber. Dynamics, Computation, and the "Edge of Chaos": A Re-Examination, 1993. Santa Fe Institute, Paper number 93-06-040. http://citeseer.nj.nec.com/4297.html.

[20] Melanie Mitchell and Charles E. Taylor. Evolutionary Computation: An Overview. *Annual Review of Ecology and Systematics*, 20:593–616, 1999. http://citeseer.nj.com/mitchell99evolution.html.

[21] C. Ofria, C. T. Brown, and C. Adami. Avida User's Manual. http://nemus.dllab.caltech.edu/avida/manual/manual.pdf and general information about Avida can be found at http://nemus.dllab.caltech.edu/avida, 1998.

[22] Eric T. Olson. The Otological Basis of Strong Artificial Life. *Artificial Life*, 3:29–39, 1997.

[23] Graham Oppy and David Dowe. The Turing Test. *The Stanford Encyclopedia of Philosophy (Summer 2003 Edition)*, 2003. Edward N. Zalta (ed). http://plato.stanford.edu/archives/sum2003/entries/turing-test/.

[24] H. H. Pattee. Artificial Life needs a Real Epistemology. 1995. Advances in Artificial Life: Third European Conference on Artificial Life, Berlin, Lecture Notes on Artificial Intelligence, pages 23-38.

[25] Thomas S. Ray. An Approach to the Synthesis of Life. *Artificial Life II*, pages 371–408, 1991.

[26] Thomas S. Ray. How I Created Life In A Virtual Universe. *Originally published in French. Submitted to Natural History Magazine, American Museum of Natural History*, 1992. http://www.isd.atr.co.jp/ ray/pubs/nathist/.

[27] Thomas S. Ray. A Proposal to Create a Network-wide Biodiversity Reserve for Digital Organisms. Technical Report TR-H-133, Advanced Telecommunications Research Institute International (ATR), 1995. http://www.his.atr.co.jp/ ray/pubs/reserves/index.html.

[28] Thomas S. Ray. Frontiers of Life, Volume 1: The Origins of Life. *Artificial Life*, pages 107–124, 2001. Academic Press. Editors: Renato Dulbecco, David Baltimore, Francios Jacob, Rita Levi-Montalcini.

[29] Thomas S. Ray. Overview of Tierra at ATR. *Technical Information, No. 15, Technologies for Software Evolutionary Systems*, 2001. Advanced

Telecommunications Research Institute International, ATR-HIP, Kyoto, Japan. http://www.isd.atr.co.jp/ ray/pubs/overview/Overview.doc.

[30] Martin Rees. Is There Life Beyond Earth? *New Scientist*, 12 July:24–27, 2003.

[31] Pritika Sanghi and David Dowe. A Computer Program Capable of Passing I.Q. Tests. 2003. Proc. Joint International Conference on Cognitive Science, Sydney, Australia, 13-17 July, 2003. (The program can be tried at http://www-personal.monash.edu.au/ psan5).

[32] Peter W. Shor. Quantum Computing. ICM I Extra Volume, Proceedings Paper, pages 467-486. http://citeseer.nj.nex.com/shor98quantum.html, 1998.

[33] Tadeusz Szuba. A Formal Definition of Collective Intelligence and its IQ Measure. *Future Generation Computing Journal*, 1999. http://ipdps.eece.unm.edu/1999/biosp3/szuba.pdf.

[34] Tadeusz Szuba and Mohammed Almulla. Was Collective Intelligence Before Life on Earth? *Lecture Notes in Computer Science*, 1800:586–??, 2000. http://citeseer.nj.nec.com/szuba00was.html.

[35] A. M. Turing. Computing Machinery and Intelligence. *Mind*, 59:433–460, 1950. http://www.loebner.net/Prizef/TuringArticle.html.

[36] Gabrielle Walker. Ice Magic. *New Scientist*, 12 April 2003:30–34, 2003. http://www-personal.monash.edu.au/ stjac1/honours.html.

[37] Scott A. Wallace and John E. Laird. Intelligence and Behavioral Boundaries. 2002. NIST Workshop on Performance Metrics for Intelligent Systems (PerMIS 2002). Gaithersburg, MD, Editors: E.R. Messina and A.M. Meystel, NIST Special Publication 990. http://citeseer.nj.nec.com/528442.html.