Artificial Life and Machine Consciousness

A. L. Nelson
Androtics LLC, Tucson AZ, USA
alnelson@ieee.org

Abstract

Over the last several decades research efforts have explored various forms of artificial life and embodied artificial life as methods for developing autonomous agents. Such approaches, although a part of the AI canon, are rarely used in research aimed at creating artificial general intelligence. This paper explores the prospects of using *in silico* artificial evolution to develop machine consciousness, or strong AI. It is possible that artificial evolution and situated self-organizing agents could become viable tools for studying machine consciousness, but there are several issues that must be overcome. One problem is the use of exogenous selection methods to drive artificial evolutionary processes. A second problem relates to agent representation that is inconsistent with the environment in which the agents are situated. These issues limit the potential for open-ended evolution and fine-grained fitting of agents to environment, which are likely to be important for the eventual development of situated artificial consciousness.

Introduction

Some authors have advocated for the use of artificial life (ALife) and evolutionary learning as a primary component of strong AI (Parisi 1997, Looks and Goertzel 2006). This paper discusses the long term prospects of using ALife to create, or at least shed light upon the question of, machine consciousness and artificial general intelligence (AGI). Largely due to automatic design aspects, machine consciousness may (at some point) be achieved by a refined version of ALife even without a working understanding of consciousness. Further, future versions of ALife may be able to side-step both the symbol grounding problem (Harnad 1990) that plagued early attempts at strong AI (Newell and Simon 1972) as well as the frame problem (Dennett 2006, Wheeler 2008, Dreyfus 2007) that has been a more recent difficulty in situated AI (Froese and Ziemke 2009).

This paper is organized as follows: First, some common terminology is presented, followed by a brief discussion of AI's failure to produce AGI or machine consciousness. The second section of the paper charts the development of AI methodologies that led up to ALife in its current form(s). After these preliminaries, an idealized (but computationally intractable) ALife paradigm is laid out, followed by a discussion of modifications to make this paradigm more feasible while preserving aspects that support the open-ended evolution of complexity (Standish, 2003).

AGI, Strong AI, and Consciousness

The field of AGI has, to a degree, subsumed the notion of strong AI (Legg and Hutter 2007). But AGI is not exactly synonymous with strong AI. A system possessing strong AI would have consciousness, not just mindless intelligence (Searle 1980). AGI research is ultimately aimed at building systems that surpass human intellectual abilities, but AGI researchers might not necessarily make claims as to whether such systems would have consciousness.

The terms consciousness and phenomenal states are used in this paper to refer to first-person experiential states. See Reggia (2013), Wheeler (2008), Dreyfus (2007), Loar (1990) for discussions of these and related terms in the context of AI.

In a general sense, consciousness might be thought of as the first-person internal experience of mind. The previous statement points to what consciousness is, but is not a true definition in the scientific sense. The experience of conscious mind is a first-person phenomenon and may not have a third-person objective manifestation that is accessible to scientific verification (Loar 1990, Arrabales, Ledezma, and Sanchis 2010).

Why don't we have strong AI?

There is no clear consensus among AI practitioners and observers as to exactly why past endeavors in all their various forms have not spawned AGI, much less machine consciousness. This failure may in part be due to fundamental flaws in methodology or it may due to the fact that few of AI's paradigms have ever been tested at full scale.

Several authors contend that the reasons for failure are related to the lack of working definitions or adequate understandings of phenomenal states, agency and consciousness (Dreyfus 2007, Menant 2012, Shani 2013, Reggia 2013). Efforts to develop a theory of phenomenal states are on-going, but there remains a lack of anything approaching a consensus (Arrabales, Ledezma, and Sanchis 2010, Koch and Tsuchiya 2011). Reggia (2013) states "no approach to artificial consciousness has presented a compelling demonstration of instantiated consciousness in a machine, or even clear evidence that instantiated machine consciousness will eventually be possible." In this vein Menant (2012) similarly concludes that we don't understand basic aspects of consciousness and this limits our ability to investigate strong AI.

Another school of thought addressing why AGI has not come to pass is that AI practitioners simply have not put in a sufficiently coordinated concerted effort. The contention is that a very large and explicitly coordinated Manhattan Project-like effort is needed to create AGI (Brunette, Flemmer, and Flemmer 2009). IBM's projects Deep Blue (Hsu 2002) and Watson (Ferrucci et al. 2010) lend some credence to this view. Watson in particular required something of a mini-Manhattan Project: three years and 20 core developers and scientists working in a closely coordinated group. These systems are impressive, but there is no consensus that either might scale up to make the general problem-solvers envisioned by AGI.

The Evolution of AI and ALife

In order to provide context, and to relate ALife research to AGI and machine consciousness, this section briefly summarizes important developments in AI, AGI and ALife.

Symbolic AI

As chronicled by Buchanan (2005), Newell (1982), Reggia (2013) and others, in the 1950's and 1960's, much of AI was focused on symbolic structured information systems, as well as various axiomatic and semi-formal methods of logical reductionism championed by Newell, Simon, McCarthy, Minsky, Shannon and others. In this paper the term symbolic AI is used to refer to this overall paradigm (Dreyfus 2007). These methods suffered from various semantic and representational problems that were much more profound than initially realized. Perhaps the most debilitating of these became known as the symbol grounding problem (Newell and Simon 1972, Harnad 1990). The symbol grounding problem asks how meaning can be associated with symbols in a computational system such as a digital computer. To a large degree this problem was never satisfactorily resolved, and led to several important theoretical and philosophical critiques of symbolic AI (Dreyfus 2007, Searle 1980, Steels and Brooks 1995, Wheeler 2008). Although all of these critiques struggled to some degree with the lack of a scientifically tractable definition of consciousness, it is generally (though not universally) accepted that high level symbolic representational systems would likely not be conscious. Even so, it is possible that systems with at least some aspects of symbolic AI may eventually gain an appreciable level of general problem-solving ability. The last decade has seen some glimmer of this in IBM's major projects Deep Blue (Hsu 2002) and Watson (Ferrucci et al. 2010), as well as programs that can perform at human levels on IQ tests (Sanghi and Dowe 2003).

An important ideological consequence of symbolic AI's difficulties and subsequent/concurrent critique was the rise of the view that conscious intelligence (at least in humans and some animals) is closely tied to an ongoing immediate interactive relationship to the environment (Steels and Brooks 1995, Dreyfus 2007, Wheeler 2008). The view was that conscious intelligent agents must be embodied, contextualized, or otherwise situated in an environment, in addition to being possessed of a tight coupling between sensors, actuators and control. This ideology received the moniker embodied AI (Steels and Brooks 1995).

Embodied AI

Although embodied AI highlighted the importance of considering a situated intelligent agent's close relationship to its environment, the control aspect of initial systems developed in the late 1980's was in some ways the weak link. Typically, these used hierarchical control architectures that were fairly rigid and hand-tuned. Although agents were environmentally situated, they had very little ability to learn or adapt (Steels and Brooks 1995). The more intrinsic aspects of control, such as fine-grained context-appropriate behavior, were not fully addressed by many early embodied AI projects. Brunette, Flemmer, and Flemmer (2009) and Anderson (2003) present reviews of embodied AI.

Some researchers, including Froese and Ziemke (2009) and others, have suggested that something fundamental is missing from our understanding of how consciousness (or any level of awareness) might arise in situated agents. Brooks in particular, after a decade of embodied AI work, wrote, "I am suggesting that perhaps at this point we simply do not get it, and that there is some fundamental change necessary in our thinking..." (Brooks 1997). This highlights a more pervasive problem in AI that has become known as the frame problem (Dennett 2006, Dreyfus 2007, Wheeler 2008). The frame problem essentially asks, "How does an intelligent agent understand its current context to the degree that it can produce appropriate action?" In short, how does an agent know, moment by moment, what to do and when to do it? In some ways this inability of early embodied AI to address the frame problem influenced the onset of ALife as a significant subfield of autonomous systems research.

ALife

This subsection focuses on ALife methods in a general way. This serves as a foundation for following discussions of how ALife might be advanced toward open-ended evolution of complexity (Standish 2003), and how this might support the ultimate (and perhaps distant) goal of producing AGI and machine consciousness. More detailed surveys of ALife in its various forms are reported in numerous works including Bedau et al. (2000), Kim and Cho (2006), Nolfi and Floreano (2000), Nelson, Barlow, and Doitsidis (2009), and Dittrich, Ziegler, and Banzhaf (2001).

Computational ALife (as opposed to wet ALife or synthetic biology) can be divided into three categories: 1) situated-agent systems, 2) self-replicating evolving program systems, and 3) artificial chemistries (AChem).

In the first category, agents (either simulated or real) are evolved in a physical environment or semi-physical simulation. The agents are evolved using various forms of population-based evolutionary algorithms. The majority of such platforms use some sort of explicit fitness function to drive the evolutionary process. These fitness functions often measure aspects of specific task completion such as distance traveled, objects located, or acquisition of simulated energy, food or materials (Nelson, Barlow, and Doitsidis 2009).

In the 1990's and through the first few years of the new millennium this class of research received a large amount of attention both in simulated situated agent and in embodied evolutionary robotics work. Such work, although still being pursued, has declined in the past few years, largely due to a growing realization that explicit fitness evaluation and open-ended evolution are largely incompatible (Lehman and Stanley 2008).

Self-replicating evolving program systems form a second class of ALife research, and include platforms such as Tierra (Ray 1991) and Avida (Adami and Brown 1994). Self-replicating program systems generally define populations of code snippets that are placed in memory in a specialized virtual machine and are run in parallel, governed by a dedicated operating system associated with the virtual machine. The agents are self-contained code sections residing in core memory, and don't necessarily represent anything beyond their own pattern in the machine. In this sense, self-replicating programs might be considered to be not just simulations, but real instances of evolving systems (Ray 1991).

State-of-the-art self-replicating program systems support populations of tens of thousands of individuals and can propagate these at rates on the order of an effective generation per second (Shao and Ray 2010). These systems are used widely to investigate empirical questions about the nature of proto-life, as well as evolutionary population dynamics and theory (Fortuna et al. 2013).

Tierra (and Avida in some configurations) does not include an explicit fitness function per se, but the OS does play a supervisory role that does not reduce completely to fundamental rules governing program behavior (Dittrich, Ziegler, and Banzhaf 2001).

AChem systems have commonalities with self-replicating program platforms, but are formulated in terms more closely analogous to molecular interactions in the natural world. See Dittrich, Ziegler, and Banzhaf (2001) for a review. In their most general form, AChem systems define sets of objects and rules that govern the low-level interaction of these objects. Resulting structures arise from these low-level interactions. These systems generally have no explicit objective function that enforces cyclic replication. Like self-replicating program systems, AChem systems have also received much attention in the last decade and are used extensively as theoretical tools.

Natural and Artificial Evolution

This section delineates fundamental properties of nature or natural law believed to enable or support the abiogenesis and further development of self-replicators capable of open-ended complexification. This set of properties is used as a guide or set of design criteria in the final sections of this paper.

Natural evolution, and the universe in general, generates exploration and open-ended search of the configuration space defined by physical law. Nature uses no explicit selection criteria per se, and all effective evolutionary pressures are consistent with the fundamental laws that govern our universe (i.e., fitness is a fully endogenous property). Selection, as such, although appearing from our distal, observers' point of view to be driving the improvement of life in its corporeal form, is from another very low-level point of view simply a subtle form of random walk through the configuration space defined by physical law. The natural processes that led to life on Earth had no design strategies beyond stochastic search, at least initially, and so far as we know, no external agent actively monitors and alters the course of evolution. In terms of physical law, there is no distinction between living and non-living material: all is constructed according to one consistent set of fundamental laws (Lange 1996).

To summarize then: nature employs no explicit fitness function, has no supervisory agent, and all aspects of the environment (and any agents it might contain) are specified using one set of physical laws.

Endogenous Artificial Life

This section begins with the definition of an idealized hypothetical ALife environment in which situated intelligent agents might (or might not) develop. Then by degrees, this will be made less ideal, but more computationally tractable. Through this process, our overall goal is to preserve the putative properties of nature that allow for open-ended complexification.

Mirroring nature, at least in a reductionist way, a hypothetical (admittedly profoundly computationally intractable) ALife research environment that would be suitable to explore the origins of life and situated intelligence on a grand scale might be employed as follows: experimenters formulate a set of elemental structures and interaction rules (analogous to fundamental physical law), and then allow this set of elements and rules to iterate over time in a computational space-like region. After the system is allowed to run for a sufficient amount of internal time, experimenters observe what has developed, modify their sets of rules, and run the system again. Any self-replicators arising in such a system would have fully endogenous environmentally consistent representation and (depending upon the set of rules being iterated) be potentially self-complexifying. But of course, actually simulating a sufficiently large space-like region at a quantum-like level is out of the question. Lloyd (2002) estimates that the upper limit on the computational capacity of a liter of physical matter is on the order of 10^{50} operations/second. Simulating this amount of material at the quantum level might require a comparable amount of computational power. Also see Grand (2012) for a recent discussion of the feasibility of simulating bulk material.

How could the fundamental properties of such a system be retained while bringing it into the range of levels of computing power that might at some point be available? The previous section proposed that a system capable of supporting open-ended development of complexity should meet the following criteria gleaned from nature: 1) the system should be free of any exogenous, explicit fitness criteria; 2) there should be no supervisory agent that treats putatively living parts of the environment differently than nonliving parts; and 3) agents should be represented fully consistently within the defining elemental rules that specify the overall environment. The goal here is to retain these general properties, while reducing the overall computing requirements to the bare minimum. Recent theoretical and research efforts have begun to address questions related to these properties, and some of these works are cited in the following discussion. In this way something of an overall trajectory along which ALife is already proceeding is sketched.

To a degree, the first and second of these properties can be achieved by omission.

AChem-like systems might potentially be tuned to fulfill all of these criteria (Faulconbridge et al. 2012). However, the computing power required to support a simulation capable of evolving agents from this basic level is likely to remain out of reach indefinitely. Addressing this while maintaining our desired criteria will require a considerable amount of finessing on two fronts. Firstly, the way in which the defining elemental rules are implemented must be dramatically altered from the basic AChem-inspired form in order to approach computational tractability. Secondly, it is probably necessary to seed environments with proto-life agents (Hickinbotham et al. 2010, Damer et al. 2010). It is likely that choosing a set of elemental defining rules that will readily give rise to lifelike agents is a much more difficult task than creating seed agents that can exist in a given environment.

There are several modifications that might allow for a considerable speed-up in the simulation of a set of rules defining a "physics" (Grand 2012). One example of these would be to compress many aspects of the environment. Still, even with the most optimistic curtailment of physical laws, it is likely that this alone would not bring us within a reasonable experimental time frame. Grand (2012) estimates that even with a set of very generous fixes, it would take in excess of 1000 years for a single instance/simulation of a set of laws to reach the abiogenesis stage (if it were to occur), and possibly a very large number of attempts to refine our set of laws.

It will be necessary to consider additional fixes that go beyond streamlining the defining elemental rules/laws of our system. One potential fix is to include many aggregated physical properties into macro-scale objects. Rules that simulate macro-scale objects as semi-elemental forms might be considered quasi-magical in relation to the elemental rules and have the potential to distort the environment outside the bounds of consistency detailed above, and must be implemented with care. One possibility regarding how to make quasi-magical rules consistent with a set of elemental rules is to require quasi-magical rules to stay within the energy confines of the simulation. Potentially, measures of complexity in ALife as studied in Lehman and Stanley (2008) could be used to modulate the degree of computational granularity in conjunction with levels of abstraction (Nellis and Stepney 2010).

Even with these modifications it will be necessary to seed the environments with proto-agents. With supporting modifications to the defining elemental rules governing interactions, it may be possible to design bodies for the seed agents, although this is currently beyond the state of the art (Ducharme, Egli, and Legault 2012). Recently, using Stringmol, an AChem system with some aspects of self-replicating programs, Hickinbotham et al. (2010) seeded the environment with a hand-designed molecule that is capable of copying any other molecule it binds to.

This then led to the evolution of a great variety of complex co-replicating molecules.

Seed agents require rudimentary sensor-control-motor loops constructed using only the defining elemental rules of the environments. This is likely to be extremely difficult within a system with simple AChem-like rules. Dualism, in which agent minds are represented externally to the environment, would violate consistent agent-environment representation. One way to address this might be to include fundamental elements/particles in the environment that are essentially functional computational units that can be coupled together into connectionist structures. Recently, something close to endogenous encoding of a rudimentary control function has been achieved in a few works. For example, Joachimczak et al. (2012) demonstrated the evolution of locomotion ability in populations of multicellular agents with genomes that specified morphology and control at the cellular level. Global or whole-agent control arose as an emergent property.

In summary, a computational ALife environment might be achieved by implementing: 1) streamlined AChem-like rules governing interaction of elements in the environment, augmented by 2) a set of rules for aggregated macro-scale objects, variable granularity based on complexity, and fundamental elements that can readily be formed into connectionist information processing and sensing constructs; and 3) initial seed populations of hand-designed self-reproducing proto-agents.

Discussion and Closing Remarks

The general methodology outlined in the previous section represents something akin to a long-term research endeavor that is already underway within the ALife research community (Shao and Ray 2010, Hickinbotham et al. 2010, Nellis and Stepney 2010, Joachimczak et al. 2012, Faulconbridge et al. 2012, Lehman and Stanley 2008) and others.

If agents in ALife environments are able to continually detect and smoothly respond to their environment (or environmental context), and if this ability is in fact one possible basis for primitive consciousness, then ALife may indeed side-step (without necessarily understanding) some of the difficulties encountered when explicitly building/programming machines with intelligence.

This paper has suggested that combining features from different areas of ALife research may produce a system that generates both open-ended development and continued complexification. Work over the last two decades has demonstrated that some level of automatic adaptation of agents to environments in situated-agent systems is indeed possible. Evolved agents are capable of performing situated behaviors of some complexity. Self-evolving program systems have demonstrated some elements of open-ended evolution, although the agents evolved are too simple to be considered complex life. AChem systems do address concerns of exogenous fitness and of inconsistent representation of evolving agents. Here though, seeding seems necessary to jumpstart the rise of complex agents, and the design of such seed agents is currently beyond the state of the art.

Philosophically speaking, there are likely to be fundamental problems in determining whether resulting agents have phenomenal consciousness even if they appear outwardly to have features we associate with consciousness. Observations of such features in agents arising from ALife simulations, though still not providing conclusive proof of artificial phenomenal states, might bear more weight than they would in super-intelligent human-designed AGIs. This is because such AGIs might have been explicitly or implicitly programmed to simulate affect, while agents created in a self-contained evolutionary simulation would have evolved such behavior independently.

Acknowledgments

The author would like to thank Brenae L. Bailey and Julia B. Nelson for many related discussions and assistance in preparing this manuscript.

References

Adami, C., and Brown, C. T. 1994. Evolutionary Learning in the Artificial Life System “Avida”. In Brooks, R. A. and Maes, P., eds. *Proceedings of Artificial Life IV*, 377–381. MIT Press.

Anderson, M. L. 2003. Embodied Cognition: A Field Guide. *Artificial Intelligence* 149(1): 91-130.

Arrabales, R.; Ledezma, A.; and Sanchis, A. 2010. ConsScale: A Pragmatic Scale for Measuring the Level of Consciousness in Artificial Agents. *Journal of Consciousness Studies* 17(3-4): 131-164.

Bedau, M. A.; McCaskill, J. S.; Packard, N. H.; Rasmussen, S.; Adami, C.; Green, D. G.; Ikegami, T.; Kaneko, K.; and Ray, T. S. 2000. Open Problems in Artificial Life. *Artificial Life* 6(4): 363-376.

Brooks, R. A. 1997. From Earwigs to Humans. *Robotics and Autonomous Systems* 20(2): 291-304.

Brunette, E. S.; Flemmer, R. C.; and Flemmer, C. L. 2009. A Review of Artificial Intelligence. In *Proceedings of the Fourth International Conference on Autonomous Robots and Agents*, 385-392.

Buchanan, B. G. 2005. A (Very) Brief History of Artificial Intelligence. *AI Magazine* 26(4): 53-60.

Damer, B.; Newman, P.; Gordon, R.; Barbalet, T.; Deamer, D.W.; and Norkus, R. 2010. The EvoGrid: A Framework for Distributed Artificial Chemistry Cameo Simulations Supporting Computational Origins of Life Endeavors. In *Proceedings of the 12th International Conference on the Synthesis and Simulation of Living Systems*, Odense, Denmark.

Dennett, D. C. 2006. Cognitive Wheels: The Frame Problem of AI. In Bermudez, J. L., ed. *Philosophy of Psychology*: *Contemporary Readings*, 433-454. New York: Routledge.

Dittrich, P.; Ziegler, J.; and Banzhaf, W. 2001. Artificial Chemistries – A Review. *Artificial Life*, 7(3): 225–275.

Dreyfus, H. L. 2007. Why Heideggerian AI Failed and How Fixing It Would Require Making It More Heideggerian. *Philosophical Psychology* 20(2): 247-268.

Ducharme, V.; Egli, R.; and Legault, C. Y. 2012. Energy-Based Artificial Chemistry Simulator. In *Artificial Life 13: Proceedings of the Thirteenth International Conference on the Synthesis and Simulation of Living Systems*. East Lansing, MI.

Faulconbridge, A.; Stepney, S.; Miller, J. F.; and Caves L. 2010. RBN-World: The Hunt for a Rich AChem. In *Proceedings of the 12th International Conference on the Synthesis and Simulation of Living Systems*, Odense, Denmark.

Ferrucci, D.; Brown, E.; Chu-Carroll, J.; Fan, J.; Gondek, D.; Kalyanpur, A. A.; et al. 2010. Building Watson: An Overview of the DeepQA Project. *AI Magazine* 31(3): 59-79.

Fortuna, M. A.; Zaman, L.; Wagner, A. P.; and Ofria, C. 2013. Evolving Digital Ecological Networks. *PLOS Computational Biology* 9(3): e1002928.

Froese, T., and Ziemke, T. 2009. Enactive Artificial Intelligence: Investigating the Systemic Organization of Life and Mind. *Artificial Intelligence* 173(3): 466-500.

Grand, S. 2012. Playing God: The Historical Motivations of Artificial Life. In *Origin(s) of Design in Nature*, 733-748. Springer Netherlands.

Harnad, S. 1990. The Symbol Grounding Problem. *Physica D*, 42: 335–346.

Hickinbotham, S.; Clark, E.; Stepney, S.; Clarke, T.; Nellis, A.; Pay, M.; and Young, P. 2010. Diversity from a Monoculture: Effects of Mutation-On-Copy in a String-Based Artificial Chemistry. In *Proceedings of the 12th International Conference on the Synthesis and Simulation of Living Systems*. Odense, Denmark.

Hsu, F.-H. 2002. *Behind Deep Blue: Building the Computer That Defeated the World Chess Champion*. Princeton, NJ.: Princeton University Press.

Joachimczak, M.; Kowaliw, T.; Doursat, R.; and Wrobel, B. 2012. Brainless Bodies: Controlling the Development and Behavior of Multicellular Animats by Gene Regulation and Diffusive Signals. In *Artificial Life 13: Proceedings of the Thirteenth International Conference on the Synthesis and Simulation of Living Systems*. East Lansing, MI.

Koch, C., and Tsuchiya, N. 2011. A Test for Consciousness. *Scientific American* 304(6): 44-47.

Kim, K.-J., and Cho, S.-B. 2006. A Comprehensive Overview of the Applications of Artificial Life. *Artificial Life* 12(1): 153-182.

Lange, M. 1996. Life, "Artificial Life," and Scientific Explanation. *Philosophy of Science* 63(2): 225-244.

Legg, S., and Hutter, M. 2007. A Collection of Definitions of Intelligence. *Frontiers in Artificial Intelligence and Applications* 157: 17-24.

Lehman, J., and Stanley, K. O. 2008. Exploiting Open-Endedness to Solve Problems through the Search for Novelty. *Artificial Life* 11: 329.

Lloyd, S. 2002. Computational Capacity of the Universe. *Physical Review Letters* 88(23): 237901.

Loar, B. 1990. Phenomenal States. *Philosophical Perspectives* 4: 81-108.

Looks, M., and Goertzel, B. 2006. Mixing Cognitive Science Concepts with Computer Science Algorithms and Data Structures: An Integrative Approach to Strong AI. *AAAI Spring Symposium Series*. Palo Alto, CA: AAAI Press.

Menant, C. 2012. Turing Test, Chinese Room Argument, Symbol Grounding Problem. Meanings in Artificial Agents. In *Proceedings of the AISB/IACAP World Congress. Natural Computing/Unconventional Computing and Its Philosophical Significance*, 51-56. Birmingham, UK.: The Society for the Study of Artificial Intelligence and Simulation of Behaviour.

Nelson, A. L.; Barlow, G. J.; and Doitsidis, L. 2009. Fitness Functions in Evolutionary Robotics: A Survey and Analysis. *Robotics and Autonomous Systems* 57(4): 345-370.

Newell, A., and Simon, H. A. 1972. *Human Problem Solving*. Englewood Cliffs, NJ: Prentice-Hall.

Newell, A. 1982. The Knowledge Level. *Artificial Intelligence* 18(1): 87-127.

Nellis, A., and Stepney, S. 2010. Automatically Moving Between Levels in Artificial Chemistries. In *Proceedings of the 12th International Conference on the Synthesis and Simulation of Living Systems*, Odense, Denmark.

Nolfi, S., and Floreano, D. 2000. *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. Cambridge, MA: MIT Press.

Parisi, D. 1997. Artificial Life and Higher Level Cognition. *Brain and Cognition* 34.1: 160-184.

Ray, T. S. 1991. An Approach to the Synthesis of Life. *Artificial Life II* 10: 371-408.

Reggia, J. A. 2013. The Rise of Machine Consciousness: Studying Consciousness with Computational Models. *Neural Networks* 44: 112:131.

Sanghi, P., and Dowe, D. L. 2003. A Computer Program Capable of Passing IQ Tests. *Proceedings of the Fourth ICCS International Conference on Cognitive Science*.

Searle, J. R. 1980. Minds, Brains, and Programs. *Behavioral and Brain Sciences* 3(3): 417-457.

Shani, I. 2013. Setting the Bar for Cognitive Agency: Or, How Minimally Autonomous Can an Autonomous Agent Be? *New Ideas in Psychology* 31(2): 151-165.

Shao, J., and Ray, T. S. 2010. Maintenance of Species Diversity by Predation in the Tierra System. In *Twelfth International Conference on the Synthesis and Simulation of Living Systems (ALIFE)*, 533-540.

Standish, R. K. 2003. Open-ended Artificial Evolution. *International Journal of Computational Intelligence and Applications* 3(02), 167-175.

Steels, L., and Brooks, R. A. eds. 1995. *The Artificial Life Route to Artificial Intelligence: Building Embodied, Situated Agents.* Psychology Press.

Wheeler, M. 2008. Cognition in Context: Phenomenology, Situated Robotics and the Frame Problem. *International Journal of Philosophical Studies* 16(3): 323-349.