

Academic Literature Review Final Draft
Noel Weichbrodt
200311.20

Post-AI: Evidence of a Paradigm-Shift in the Program of Artificial Intelligence

Abstract

Traditionally, the field of Artificial Intelligence has been dominated by the philosophy and program of Strong AI. Strong AI, advocated notoriously by Ray Kurzweil and others, maintains that human-level intelligence may be replicated by a combination of advanced algorithms and increasing computing capacity, and directs the AI research program to create a human-level intelligence. That task has met with much unexpected and unsolvable difficulty, and there are good reasons to think that Strong AI in both view and action is completely wrong. In place of that, a weaker view of AI -- that we should try to create animal-level intelligence -- may have fecundity and congruency that Strong AI does not. Starting from David Hume and moving on to modern literature by Douglas Hofstadter and other, this paper surveys promising directions in research that this new sort of artificial animal programs may take.

1.0.0 Strong Artificial Intelligence (AI), the publicly popular stance in AI, maintains that human intelligence is replicable computationally, and AI's project is to create a human-level intelligence. Weak AI holds that the creation of a human-level intelligence is impossible and/or unfeasible.

The image is almost archetypal: a metal man, with a head that glints and pulses in hues of LED, spindly arms and legs, and an clean and humming electronic brain that is more efficient than our wet, messy, biological unit. From the first science-fiction stories of “robots” to the 1997 victory of the Deep Blue computer over the World Chess Champion [Hsu 2002], tantalizing hints that humans are on the edge of replacing themselves with machines of their own creation crawl thick through the cultural air.

The fulfillment of this image is the goal of what is known as the “Strong” view of Artificial Intelligence (Strong AI, see Appendix for further elucidation on the taxonomy of Artificial Animal research programs). Strong AI has quite a few backers in the areas of philosophy, science, and media. The inventor of the modern synthesizer keyboard and prolific inventor, Ray Kurzweil, is the exemplar of this view. He believes that Strong AI will, within the next 30 years, create a human-level intelligence, and after that we will enter the post-human evolutionary era [Kurzweil 1999]. The philosopher Marvin Minsky believes that human intelligence is simply a complex, but replicable, biological process that can be modeled algorithmically. [Minsky unpublished] Transhumanist Kevin Warwick is attempting to enhance and eventually replace his biological body with a synthetic one. [Asohan 2003]. And the creators of the blockbuster *Matrix* films imagine a future where machines rule earth, and humans are forced under the ground to plot a rebellion against their newly-fashioned oppressors.

There is one problem with this image, though: it will not happen. Of course this point is under contention, but the view of the Weak AI movement is that the replication of a human-level intelligence is a hopeless task, and that research should be conducted in more feasible and fruitful realms. Research into creating artificial

insects, mice, and even dogs has gained popularity in recent years. AI like this seems to be near to grasping the key concepts of animal-level intelligence. Meanwhile, mainstream Strong AI research has bogged down in issues like the role of the symbol in human thought, and the lack of progress based on the assumption of the emergence of consciousness out of complexity.

1.1.0 Strong AI's Program: by a combination of clever programming and large computational resources, create a machine that "thinks".

That intelligence and consciousness can arise out of complex computation is perhaps the central assumption that Strong AI makes. There are several parts to creating complex computations. First, there must be complex algorithms. This has been the most difficult area of AI research, and there seems to be much work still needed to create sufficiently complex algorithms. Second, computational ability must increase—from to gigaflops to teraflops to petaflops. Third, storage capacity and accessibility must continue to increase in space and decrease in access time.

1.1.1 The Project: create smart algorithms, ride Moore's law, keep increasing raw knowledge of computers.

Strong AI's chief challenge is to create smart algorithms, algorithms that mimic human thought in form or by function, or perhaps both ways. Strong AI theorists often hold that the chief aspect of human thought consists of pattern recognition. [Minsky unpublished] Douglas Hofstadter and others have been working on creating computer programs that can recognize numerical patterns, create variations on a theme, and draw analogies between different concepts. [Hofstadter 1994]

Moore's Law, the famous formulation given by Intel cofounder Andy Moore that processors will double in speed and halve in cost every eighteen months, has been the direct cause for the current surge in technology growth and optimism in the future success of Strong AI. Moore's Law has caused Kurzweil to calculate that the processor will exceed the human brain in raw computational ability by 2030 [Kurzweil 1999, i], which he sees as the important landmark, the singularity beyond which humanity will never be the same.

Combine Moore's Law with the even faster growth in storage, and you have an

interesting proposition: there could possibly be a computer in 30 years that will have over 90% of the total amount of information ever known by man, contained on the internet, and at the command of some incredibly smart and fast algorithms. Will the combination of processing power and information-access combine to replicate human intelligence?

1.2.0 Weak AI's Program: find other intelligence and uses for AI.

Perhaps not. Intellectual dissent with the Strong AI program has gone back a far ways. Hubert Dreyfus in the early 1970s opined as to *What Machines Can't Do* [Dreyfus 1972], while Joseph Weizenbaum fundamentally disagreed with those who declared his 1961 ELIZA program to have passed the Turing Test [Weizenbaum 1966] . Lately, it seems that others are joining in the recognition that Strong AI is dead in the water, unable to make any headway against the intractable problems inherent in the replication of human intelligence. What does Weak AI propose as a good project in the place of humans?

The replacement ideas vary from Artificial Life to Artificial Pets [See Appendix]. But they hold in common the idea that its entirely possible to create an artificial intelligence that mimics animals. Animal-level intelligence has a broad scope that ranges from cockroaches to dolphins. But the idea is that we can start by creating cockroaches, and eventually move up to a fully-dog-like Sony AIBO. Along the way, there are fascinating and fruitful areas like ant colonies and the eternally-replaceable, never-dying, no-animal-testing-here lab rat.

1.2.1 There is a strong possibility that an animal-level intelligence can be created.

This idea is not new. In fact, the popular 18th century philosopher David Hume's *Treatise of Human Nature* notes that animals are quite basic and predictable in their function [Hume 1888, I.I.XII]. The implication, at least to any good modernist, is that animals are reproducible in a computational, algorithmic manner.

Children's toys has been a fruitful, popular, and too-early application of this. In the early 1980s, there were attempts to create children's toys, such as "Teddy Ruxpin", that would act like animals, or at least a fantasy version of animals. Then came the "digital pets" such as Tamagotchi. These were not truly artificial animals,

forcing children to interact with them in their own artificial world rather than existing autonomously in our own world. Sony brought the AIBO artificial dog out in 1998, and though it fails to come close to fully capturing the behavior of a real dog, it does provide a good instruction of where to proceed from here.

Obviously, the dolphin would be harder to replicate computationally than the cockroach, but both lie inside the scope of possibility. Though, like Strong AI, there has not yet been a successful cybernetic replacement for any organism, there seems to be greater progress and greater likelihood of success in creating an artificial ant than an artificial human. Not only is the technical achievement more possible, but the success of ALife or Artificial Pets sidestep the possibly terminal philosophical, ethical, and spiritual issues wrapped up in the attempt to create a human artificial intelligence.

1.2.2 The pursuit of a Weak AI program will give more results than Strong AI, and it congrues more with the way nature is and the way we create.

Earlier, there was talk of the idea of fecundity. Fecundity is a familiar concept in philosophy of science, referring to a paradigm or theories' fruitfulness, its ability to generate new solutions, or approaches, to existing problems, preferably problems in a completely unrelated field. Though Strong AI has had some fecundity, especially in the areas of pattern recognition and expert systems, it still has not really impacted the technology world at large. Work in ALife, though younger than Strong AI, has already lead to interesting new ideas in unrelated fields. For example, Mitchel Resnick's work into stimulating microworlds has found new approaches to educational theory and practice, and in the development of the modern toy LEGO Mindstorms [Resnick 1994].

2.0.0 Strong AI is wrong.

There are a myriad of issues wrapped up in the claims of Strong AI. Fundamentally, Strong AI is a belief system, not purely a research program [Lanier 2000]. In that, note the inherent ideologies of reductionism, adaptationism, emergentism, and materialism. The mantra of Minsky [Demski 2002] is "the mind is a computer made of meat." Strong AI has idealist beliefs regarding the ability of

humans to surmount their own problems regarding complexity of computer engineering [Brooks 1995], and that machines can at some point start a program of efficient, serious, sustained, and intelligent self-evolution. These philosophical beliefs may or may not be true, but they are presuppositions that underlie the proclamations of Kurzweil *et al.* If you find fault with them in this deconstruction, then there is reason to doubt the validity of the whole enterprise.

2.1.0 Strong AI has been unfruitful, with great claims coming and going, and not much being accomplished to meet its stated goals.

The promissory materialism [Demski 2002] that keeps swearing that stunning results are just around the corner has worn thin after 30 years. There are two commonplaces [Lanier 2000] at this stage of the argument against Strong AI. Both are specific instances of failed attempts to create a serious artificial intelligence that is capable of substantial contributions to human work.

The first is the failure of the aptly-named “expert systems”, now re-branded as “intelligent agents”. These systems were supposed to be able to give the advice of an expert when presented with normal questions, which vast databases of relevant information and one-hundred percent accuracy. An example would be a medical expert system, able to diagnose patients like a human doctor. To date, there have been no wholly successful expert systems in this, or in any other field.

The second failure is perhaps more well-known. Microsoft, which seems to crop up in every counterexample to techno-utopianism, tried to incorporate early successes in Strong AI into its office software suite and operating system. Remember “Microsoft Bob” and the dancing paperclip from Office 98? Surely, these are worst-case examples, but after thirty years of sustained Strong AI work, we are still incapable of incorporating even the most basic of the research algorithms into useful commercial products.

2.1.1 This is the fifth decade of serious work on Strong AI, and the number of known problems has only increased.

Indeed, though a number of basic approaches to creating a human-level intelligence have been explored, none have really developed at the hoped-for rate.

Neural nets are fun and useful, but they are still incapable of acting anything like a human brain. Pattern-recognition software has progressed in recognizing the human voice and optical characters, but both tasks are still judged as not ready for mass application after twenty years of sustained (and economically demanded-for) research. Witness that this paper was typed in using a keyboard (forced by the author's predilection to listen to music such as Radiohead's *OK Computer* while writing, as current voice-recognition software does not function with such background sounds), and the bibliography entries keyed in by hand. As Strong AI proponent Rodney Brooks says,

“...Both fields have been labeled as failures for not having lived up to grandiose promises. At the heart of this disappointment lies the fact that neither AI nor Alife has produced artefacts that could be confused with a living organism for more than an instant. AI just does not seem as present or aware as even a simple animal and Alife cannot match the complexities of the simplest forms of life.” [Brooks 2000, 409]

More than these failures, though, are the ever-increasing number of questions. Sure, there are modelable neural nets, but the brains of even simple animals are still mysteriously unable to be modeled by them. What's going on inside that head? What are the brain's processes that enable human-level pattern recognition? Where do the distinctives of humans spring from?

2.1.2 The distinctiveness of humans has not even been pinpointed yet—is it intelligence, emotions, a soul, a spirit, freewill, imago dei, or creativity?—let alone duplicated on a computer.

The trick is, the distinctives of humans are not even known, let alone replicated computationally. Depending on your faith, free will, creativity, pattern recognition, noble emotions, abstraction, problem-solving, the soul, or the imago dei are debatable-correct answers. Like a multiple-choice test that you failed to study for, there is a large amount of seemingly valid options, and the truth may be nonexistent, not listed, in just one, in several, or in all (note the number of options for human distinctiveness thrown around in [Kurzweil 2002]).

If the reality of what makes humans human is not known, what makes Strong AI

believe that it is possible to model computationally? The reductive materialism inherent in the assumption that this is possible strikes quite the wrong chord. Any attempt, by reductive materialists or others, is going to be dependent on their chosen metanarrative's ontological views of human nature. Fundamentally, the issue of human intelligence is subjective, and not in the domain of science.

2.2.0 Strong AI has the wrong view of what being human is. "Intelligence" is complex, and it is just one part of being human. In light of this, it is likely that Strong AI can never succeed.

Strong AI, then, can never succeed. Intelligence is incomprehensible; the characteristics of human-ness are not able to be nailed down outside of one's conceptual framework. No matter how sophisticated the algorithms, massive the storage, or speedy the calculations, there is a limit to what computers can model. Perhaps this limit may be reached in 2030, but the limit will be there, and computers will never be able to recreate human intelligence. Reaching a human-level computational ability will not by itself be able to solve all computational problems. Brooks again points out that "Even when we have computers with the same level of processing power as the human brain, they will not be able to play a good game of Go using brute-force search alone." [Brooks 2000, 410]

2.2.1 The pathological idea of intelligence: intelligence will always be just beyond what computers can do. Taken from mathematics, where a pathological sequence is one where the needed number is just one step past what can be reached.

In mathematics, a pathological sequence is characterized by the needed number being just one step past what can be reached with the sequence. Intelligence, as a summary term for the distinctive characteristics of humans, is pathological. It will always be just beyond what computers are capable of. Put negatively, we can always redefine intelligence to mean whatever it is that computers cannot currently do. So Strong AI has a constantly moving target, never staying put long enough to be hit.

2.2.2 Intelligence is still ill-defined, and that is not even the full extent of human-ness. There's the soul, the work-in-the-world, and the imago dei.

From the Christian metanarrative, intelligence compromises three things: the person, which will live eternally, the calling, both general and specific, to work-in-the-world, and the imago dei. Each of these is unique to humans, and together they form the distinctiveness of human nature.

The person is perhaps better known traditionally as the soul, that eternal part of humans which contains their distinctive traits. There is some debate as to whether souls are real, or simply placeholders for the complex sustaining action of God, so by including the whole person this is avoided. The person is born, lives, and dies, and after death finds life, either eternally with God or in Hades.

Humans have a distinct role in creation, both general and specific, spiritual and immanently physical. The general calling given to humans is to rule the earth and subdue it as God's stewards, and to glorify God and enjoy Him forever. The specific callings are given to each person as his or her life's work, and can vary from person to person, and season to season.

The imago dei, or image of God, is the mysterious stamp that God placed on humans when he created them. Some combination of our physical makeup, intelligence, relational capacity, and spiritual side are reflective of God Himself. No other creature in creation has this stamp.

3.0.0 Survey of directions in Artificial Animal research

Given all this, there are other, better directions for research than the Strong AI program. Weak AI, though lacking the hubris and grandiosity of exceeding human intelligence, still presents a number of fruitful and intriguing research areas. In spaces like complex systems simulation, adaptive systems, neural networks, and heuristic modeling there is still much work to be done. For a grand inspiration to motivate the program, look at the "Found: Artifact from the future" picture on the back page of the November 2003 issue of *Wired*. [*Wired* November 2003, 121] Coyle staring at you, a lovable brown mutt squats inside retail packaging. The label calls the dog the Sony AIBO 20th Anniversary Model, with a list of features like "Already housetrained!" and "Includes Security3000 module!" Research into animal intelligence will have progressed an admirably long way if that picture is accurate in twenty years.

3.1.0 Keep thinking about and researching what makes humans intelligent, and how we can program computers to act intelligently.

Rather, we should drop the reductive materialism that motivated the Strong AI program, and still keep trying to figure out what makes humans intelligent and how computers can act intelligently. The work of Weizenbaum, Kurzweil, Brooks, Minsky, and Hofstadter is valuable, even if it does not lead to the creation of an artificial human-level intelligence.

3.1.1 The quest to program intelligence will not succeed, but will fail interestingly.

The Strong AI program will fail interestingly, giving valuable lessons and leads in its death. That a research program is doomed but still useful can be seen in the quest for a grand unified theory in quantum physics. Though none has been found yet, the effort has led to a number of advances in other areas. This fecundity could be true too for programming intelligence. For example, a further understanding of human intelligence and computer intelligence will eventually lead to the creation of a useful, actually intelligent, Office paperclip assistant, though it will most likely not be first made by Microsoft.

3.1.2. Hofstadter's work in Creative Analogies and Fluid Concepts.

Hofstadter exemplifies the ideal of a more humble approach to researching human and computer intelligence. In *Creative Analogies and Fluid Concepts*, he works both negatively and positively in attempting to understand what human intelligence is and how computers can act intelligently. He negatively assaults bad instances of anthropomorphization in AI claims, and positively gives research into his ideas of what makes humans intelligent, specifically abstract pattern recognition and isomorphic creative concept play.

Negatively, Hofstadter notes that AI researchers often give computers strings of words that mean exactly nil to the machines, but by a combination of imagination and laziness are given a huge amount of complex meaning by humans [Hofstadter 1994, 166]. Such fraudulent usage of common terms gives meaningless strings of chars like “birth (woman, midwife)” an implied significance of “the computer understands that a woman gives birth to a new baby with herself and the help of a midwife”, when

really “answer-to-the-ultimate-question (life-universe-everything, 42)” has the exact same meaning to the computer.

Positively, Hofstadter researches a number of what he calls “microdomains”, small areas that begin to peck away at key areas of human intelligence. His program Copycat tries to make analogies by taking a given metaphor and twisting it to give it new words, but the same meaning. LetterSpirit, another program, creates new fonts of 26 characters that are similarly artistically inspired. These types of research, though small to begin with, will be fruitful and vital in any future use of computer intelligence.

3.2.0 Focus on creating an animal-equivalent intelligence. It is possible, and would be fruitful and profitable.

The shift in AI research from humans to animals holds further promise. Hume thought so, and recent papers in multi-agent systems, autonomous-agent systems, and autonomous virtual creatures show promising results.

3.2.1 Hume’s ideas of animal and human intelligence.

David Hume maintains that both animals and man have reason, and only differ in their reasoning process by degrees [Hume 1888, I.I.XVI]. Obviously much of this paper disagrees with that claim, but what Hume says could provide motivation, and certainly provocation, for those who do not accept the metaphysics put forth in this paper to further research into animal intelligence and how it differs from that of humans. Of course, Hume also denied causality, so take his claims with a large grain of salt, but he draws a few key features of reason that they both hold in common.

Self-preservation, the obtaining of pleasure, and the avoidance of pain are the key goals of both species. Both learn by habituation, and both rely on a combination of reason and instinct [Hume 1888, I.I.XII]. If the structure and operation of the parts are the same in animals as in men, then the cause of their actions must be the same, and so animals possess the same passions as men [Hume 1888, II.I.XII]. Everything operates on “springs and principles” regarding the passions in both men and animals, though animals do not portray the penetration or reflection that mark human reason. [Hume 1888, II.II.XII] So Hume, though quite a reductive materialist, notes that there

are modelable aspects in animals that contain similarities to human life.

3.2.2 Artificial Life, as focused on an Autonomous Virtual Creature.

One perhaps unknowing example of a research program based on Hume's speculations is given by the ALife research into autonomous virtual creatures. Burke and Blumberg in a recent paper [Burke and Blumberg 2002] examine how adding a temporal awareness to a virtual creature's system changes the behavior of the creature. The ability to map thoughts to time creates indeterminate and determinate beliefs, which are an important aspect of animal intelligence (and human intelligence too). [MacIntyre 1994, 38-40] Note here how easy and guilt-free it is to call a fairly complex, but still basic, set of algorithms running together in a program a "creature", as opposed to intelligence.

The system normally used in creating a virtual creature is behaviorally-based, and not, as with Strong AI, representationally-based. The important goal is for the creature to behave like an animal, and not necessarily to think like one. To that end, the most successful way so far has been to create a set of associations, responses to stimuli, and let the reactions guide the creature's behavior. The associations can be separated into three components: a "representation of the internal and external worlds" that "represent the world," an "action selection mechanism" that "decides what to do," and a "navigation and motor system" that "figures out how to do it" [Burke & Blumberg 2002, 327]. Burke and Blumberg add to this list the ability to remember things that have already happened, keep a list of predictions for the future and present based on the past, and draw lines of relation between the three simple temporal stages. By the addition of a causality-producing temporal model, the autonomous virtual creature is able to be trained using a clicker, just like a real lab rat.

3.2.3 Simulations of ALife macro creature behaviors

Mitchell Resnick [Resnick 1994] abstracts and explores the behaviors of amoebas, ants and termites [Resnick 1994, 59] at the group level. He picks features of their group behavior to model, and, using his programming language STAR-Turtle, experiments with simple models of their behavior and environment. While the

simulations are quite basic, the results are still interesting on both an empirical and structural level. The simple rules that govern the animals, often easily described in two or three lines of code, bring about counter-intuitively rich behaviors that mimic aspects of ALife groups. Structurally, the simulations reveal the way humans perceive animals.

3.3.0 How an animal intelligence would integrate in a Biblical ontology.

This speculation brings up a final question: what is the ontological status of a synthetic creature, an artificial animal intelligence? Do we afford it the same status as normal animals? Do we treat it like other, simpler robotic toys? Does it make a difference if the creature is purely virtual or purely physical?

3.3.1 Artificial Pets do not have the same ontological status as natural animals, but is still something we must steward wisely and discerningly.

Animals that are artificially created are not the same ontologically as the created animals. To give a thought experiment that will show this, consider your 20th Anniversary Edition Sony AIBO, Fido. Fido is in every way the same as your other, normal, mutt Sugar, except that Fido never misbehaves and is much more reliable and watchful, and lower maintenance than Sugar. Is it morally wrong to do things to Fido that would be morally wrong to do to Sugar? For example, say you got tired of Fido's earnest perfectionism, and tossed the poor dog out of your 10th story apartment window. Fido is smashed to bits, but your life is happier. Is that wrong? For further consideration, do you punish your three-year-old Alice the same when she pulls Fido's tail as when she pulls Sugar's? Is it a "slippery slope" from mistreating an artificial dog to mistreating a real dog, much like cruelty to animals is supposed to desensitize actual child abuse?

The answers to these questions, are no, no, yes, and yes, respectively. But that is a simplification, and the real world will provide examples that do not have snap answers. Artificial animals are not the same as real ones, but the difference is difficult to navigate. If the Artificial Pet project succeeds like it should in the next twenty years, these questions will be faced in our lifetime.

3.4.0 Conclusion

We have surveyed the case against the Strong AI program, and dismissed it due to its unfruitfulness, arrogance, disagreeable presuppositions, and terminally confused ideas about just what intelligence is. In its place, a program of Weak AI exploring the creation of animal-level intelligence is considered. This program does not carry the same presumptions as Strong AI, and so stands a chance of success and fruitfulness. The new program is two-pronged—continue to explore just what makes humans intelligent and how computers can model parts of that, and work to create artificial animal-level intelligence. Examples from Hofstadter and Burke & Blumberg show some hopeful results in these areas. Finally, the ontological status of such artificial creatures was explored, with the conclusion that they are not the same as real creatures, but still worthy of stewardly consideration.

Appendix: A Taxonomy of Artificial Animal Research Programs

This taxonomy exists to set out positive research programs for the various levels of life as it exists on Earth. Two positions, Weak AI and Weak ALife, are not programs as such but instead negative responses to positive programs, and define themselves solely in opposition to other programs. The divisions made flow from three considerations. The first is where and what the current programs of research into these areas divide themselves as. The second accounts for the role that perception plays in the animal's behavior. The third consideration in the divisions within this taxonomy is what roles intention and communication play in the animal. Note also that I include humans as animals here, as per [MacIntyre 1994]. The taxonomy addresses four points on the scale of life, of which MacIntyre states:

“Ourselves, dolphins, chimpanzees, dogs, bats, lizards, and spiders [exist] as at different points on a scale [of animals]...At the lower points on that scale perception plays a part in the causation of behavior, perhaps as information-affording, but not because it is reason-affording...At higher points what part perception plays in causing behavior sometimes varies with just how far it is reason-affording or is taken to be reason-affording.” [MacIntyre 1994, 59]

MacIntyre elucidates further on the role of intention and communication, “Differences in the type of intention exhibited in behavior and communicated to others are also important.” [MacIntyre 1994, 57]

Finally, this taxonomy takes the ultimate goal of each program to be the creation of a real-world synthetic animal, one that, if not mimicking the natural animal exactly, is minimally equivalent in terms of features and capabilities of the natural animal it is based on. So, for example, an end in the Artificial Pet program would be the creation of a replacement for the household dog such that the Jetsons would not be impoverished by replacing Astro with a *n*th-generation Sony AIBO. This is a slight redefinition of what the stated goal of ALife is, but one that I think is necessary for the program's future. Meatspace will always be a vital, valuable place, no matter how sophisticated the virtual world may become.

Strong Artificial Intelligence (Strong AI): Create artificial analogues to humans. These artificial animals can fulfill the pathological definition of intelligence. An example of a pathological definition is that “Only intelligent beings can make stupid mistakes.” [Slashdot.org 2003]

Weak Artificial Intelligence (Weak AI): There is no way to create artificial analogues to humans as stated above. There may be the possibility of creating lower-level animals, though, as in Artificial Pets and below.

Artificial Pets (AP): Create artificial analogues to higher-level animals: dogs, apes, dolphins, etc. These are the “Animals whose perception are in part the result of purposeful and attentive investigation and whose changing actions track in some way the true and the false.” [MacIntyre 1994, 57], who are

“...members are capable of sophisticated interaction with human beings, interactions in which the perceptions, beliefs, reasons for action, and intentions of the nonhuman participants play very much the same part as do the perceptions, beliefs, reasons for action, and intentions of the human participants.” [MacIntyre 1994, 58]

“They too inhabit a world whose salient features can have this or that significance for them. They too respond on the basis of their classifications and interpretations. They too make and correct mistakes.” [MacIntyre 1994, 60-61]

Artificial Animals (AA): Create artificial analogues to the lower-level animals: lizards, turtles, woodpeckers, etc. These animals show features of intentional intra-species communication. “[Those] types of animal for whom sense perception is no more than the reception of information without conceptual content. There is, in Heidegger’s terms, no ‘as-structure’ whatsoever.” [MacIntyre 1994, 57]

Strong Artificial Life (Strong ALife): Create artificial analogues to ants, roaches, moths, crabs, etc. Mostly in the domain of insects with primitive perceptions and communication. Boden defines ALife’s distinctive aim as “...to define simple reflex-like rules from which the more complex target-behaviour will emerge...focusing on simple processes that work bottom-up to generate order at a higher level.” [Boden 1996].

Weak Artificial Life (Weak ALife): Such artificial analogues as above are not alive in

any reasonable sense, and will only truly exist in a virtual realm.

Bibliography

- Asohan, A.: “Leading Humanity Forward”, Star-TechCentral.com, available at <http://star-techcentral.com/tech/story.asp?file=/2003/10/14/itfeature/6414580&sec=itfeature>, visited 20 November 2003.
- Boden, Margaret: *The Philosophy of Artificial Life*, Oxford University Press, 1996.
- Brooks, Frederick: *The Mythical Man-Month*, Addison-Wesley, 1995.
- Brooks, Robert.A.,
- C. Breazeal (Ferrell), R. Irie, C. Kemp, M. Marjanovic, B. Scassellati and M. Williamson: "Alternate Essences of Intelligence", *Proceedings of the Fifteenth National Conference on Artificial Intelligence (AAAI-98)*, Madison, Wisconsin, July 1998.
 - "The Relationship Between Matter and Life", *Nature*, Vol. 409, January 18, 2001.
 - "From Robot Dreams to Reality", *Nature*, Vol. 406, August 31, 2000.
- Burke, Robert, and Blumberg, Bruce: “Using an Ethologically-Inspired Model to Learn Apparent Temporal Causality for Planning in Synthetic Creatures”, *AAMAS’02*, July 15-19, 2002, Bologna, Italy.
- Crevier, Daniel: *AI: The Tumultuous History of the Search for Artificial Intelligence*, Basic Books, 1993.
- Dreyfus, Hubert: *What Computers Can’t Do: A Critique of Artificial Reason*, Harper & Row, 1972.
- Eriksson, Inger V.: “Computers as Tools”, *Ethics in Computer Age*, 1994.
- Griffin, Donald: *Animal Minds: Beyond Cognition*, University of Chicago, 2001.
- Hofstadter, Douglas and the Fluid Analogies Research Group: *Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought*, Basic Books, 1995.
- Holland, John: “Outline for a Logical Theory of Adaptive Systems”, 1961.
- Hsu, Feng-Hseing: *Behind Deep Blue*, Princeton University Press, 2002.
- Hume, David, *A Treatise of Human Nature*, Oxford, 1888.

Kurzweil, Ray:

–*The Age of Spiritual Machines*, Viking, 1999.

–Richards, Jay W., and Dembski, William A.: “Live Moderated Chat: Are We Spiritual Machines?”, *International Society for Complexity, Information, and Design* July 19, 2002. available at

<http://www.kurzweilai.net/articles/art0511.html?printable=1>

Lanier, Jason, “One-Half a Manifesto”, *Wired* 08.12, December 2000. In longer form at http://www.edge.org/3rd_culture/lanier/lanier_index.html.

MacIntyre, Alasdair: *Dependent Rational Animals: Why Human Beings Need the Virtues*, Open Court, 1999.

Minsky, Marvin: *The Emotion Machine*, unpublished manuscript, available in review form at <http://web.media.mit.edu/~minsky/>, visited 20 November 2003.

Resnick, Mitchel: *Turtles, Termites, and Traffic Jams: Explorations in Massively Parallel Microworlds*, MIT, 1994.

Slashdot.org: “Kasparov Wins Game 3 Against X3D Fritz”, available at

[http://games.slashdot.org/comments.pl?sid=86082&threshold=3&](http://games.slashdot.org/comments.pl?sid=86082&threshold=3&commentsort=0&tid=127&tid=186&mode=thread&pid=7490085#7490566)

[commentsort=0&tid=127&tid=186&mode=thread&pid=7490085#7490566](http://games.slashdot.org/comments.pl?sid=86082&threshold=3&commentsort=0&tid=127&tid=186&mode=thread&pid=7490085#7490566), visited

18 November 2003.

Tu, Xiaoyuan and Terzopoulos, Demetri: “Artificial Fishes: Physics, Locomotion, Perception, Behavior”.

Wiener, Norbert: *God and Golem, Inc.: A Comment on Certain Points where Cybernetics Impinges on Religion*, MIT, 1964.

Wired: “Found: Artifact from the Future”, *Wired* 11.11, November 2003.