

# Failure and Success in AGI Projects

Thomas S. McCabe

Rensselaer Polytechnic Institute  
Buck 1001, 1999 Burdett Ave., Troy, NY, 12180  
[pphysics141@gmail.com](mailto:pphysics141@gmail.com), [mccabt3@rpi.edu](mailto:mccabt3@rpi.edu)

## Abstract

Building an artificial general intelligence, like any other difficult project, can only be successfully completed under the right environmental conditions. In order to make substantial progress on the problem of AGI, we will need to ensure that the interests of the AGI research community favor the production of a complete, working theory of general intelligence, and simultaneously avoid the various pitfalls of large research projects which have plagued other fields over the past two thousand years.

## Cascading Failure

Over the past fifty years, the field of artificial general intelligence has, time and time again, proposed building an AGI, gotten research funding, published papers, and then failed to deliver anything close to a human-level reasoning system (Lighthill 1973, Howe 2007, Russell 1999). Although the field of narrow AI has indisputably made huge gains during this time period, this success has been largely overshadowed (Kurzweil 2005) by the failure of artificial general intelligence, which has come to dominate the public's perception of AI as a whole. It is proposed that this pattern of failure and disillusionment is *not* caused by the shortcomings of individual AGI researchers, or by any difficulty inherent to the task of building an AGI; rather, it is a systematic organizational failure, which has caused corporations and universities to misdirect their problem-solving efforts for decades. This type of failure is not unique to AGI, or even to computer science in general, and similar situations have occurred in the historical record since the invention of scientific thought over two thousand years ago. By learning from the past, and by understanding the psychological and structural failures behind the history of failed attempts, we can build AGI in a timely and efficient manner, and lay down an example that will help future generations to solve their own intractable problems.

## Misdirected efforts

It is an axiom of economics that people respond to incentives (Landsburg 1995). The most important question to ask is then: what are the incentives given to researchers in AGI and the wider field of AI, and how are they responding to those incentives? It goes without saying that there is a huge incentive to build a working, human-level reasoning system, as the resulting benefit, both to the

individual researchers and society as a whole, would undoubtedly be enormous. The profound impact that such a system would have on our society has already been discussed in great detail elsewhere (Kurzweil 2005, Yudkowsky 2008), and so it will not be covered here.

However, it should be noted that this incentive only comes into play at the end of a long, tedious process of research and development; nobody expects to win the Nobel Prize for building an AGI that is only 90% complete, let alone 5% complete. And human psychology is, in very large part, much more responsive to immediate rewards than to distant rewards; in the field of heuristics and biases, it is well known that people will overvalue near-term rewards at a rate far above the rate of inflation, a phenomenon referred to as hyperbolic discounting (Green 1999). Psychologically, it appears that the degree to which people's actions are influenced by a chain of reasoning decreases exponentially with the length of the chain (Hofstadter 1996), and long-term incentives will usually have a much more indirect relationship between actions taken now and rewards received later. Hence, we should expect that the incentive structure in AGI, or in any field, will be dominated by short-term events and short-term rewards, even though the long-term rewards are much greater in magnitude.

What are the short-term motivations of researchers in AGI, or in any other field? As a general rule, they tend to be more mundane than winning the Nobel Prize or revolutionizing human society; things like publications in top-notch journals, prestige among peers, funding for grant proposals, and tenured jobs come to mind (Latour 1992). The primary way of obtaining these is to publish papers and give talks on good, original ideas, and hence, people are motivated to produce a steady supply of new concepts and paradigms in AI journals and conference proceedings. As an example, since 1980, over a hundred thousand papers have been published on artificial neural networks alone (Zeev 2004). However, because AGI is far too complicated for any single idea to be a complete solution, these new ideas rapidly get tested, and then discarded as "failures". Because of this effect, there is a strong incentive to not spend much effort following up on previous work-publications which utilize a lot of old ideas ipso facto associate themselves with the failures of those ideas, while simultaneously reducing their own originality. The final result of this process is somewhat paradoxical- a large

number of worthwhile innovations in artificial intelligence are produced, but the original problem of creating AGI is never solved, and many of the innovators in the field eventually wind up relabeling their work to get away from the general stigma of failure (Phillips 1999).

Since this process should continue for as long as the original incentive structure remains in place, we should expect to see cases similar to this in the historical record: places where huge amounts of effort were ostensibly spent on attempts to solve a problem, over a long period of time by many different groups of scientists, and the problem remained unsolved. Although the history of such incidents is far too extensive to cover here, a few examples are presented below to help outline the general principle.

**Alchemy:** Relative to the technology of the time, the ability to transmute lead into gold would have been almost as important during the Middle Ages as AGI is to us today. But due to the circumstances prevailing in the ancient world- most notably, the lack of any central, easily accessible body of knowledge on any subject- the alchemists failed to build on previous work, the pattern of 'discover, experiment, fail, and start over from scratch' developed, and the original problem was never solved until the advent of modern science.

**Epistemology:** The study of knowledge and its acquisition traces back to the ancient Greeks, and by the 18<sup>th</sup> century, there was a huge body of literature in existence on the topic. However, two millennia of effort failed to uncover even simple ideas- Laplace's law of succession, an extremely simple rule of probability theory, wasn't discovered until the early 19<sup>th</sup> century. Humanity did, eventually, succeed in understanding inductive reasoning and the use of evidence in acquiring knowledge, but only after widely accepted mathematical formalisms were introduced (Jaynes 2003).

**Unified field theories:** Unlike the previous two cases, the quest for a unified theory of physics has always been based on mathematics and hard evidence. However, the theoretical physics community also shares most of the incentive structure that dominates AGI research. As a result, once particle physics theories went beyond the testing abilities of the experimentalists, researchers rapidly created a huge assortment of new theories, none of which are easily testable (Woit 2007), without ever resolving the fundamental question of how to unify quantum theory and general relativity.

### Alternative theories

The long history of AGI, and the large amount of attention devoted to it in the academic community and the wider world, has resulted in a number of theories about why past AGI efforts have either failed, or produced results far less significant than those originally promised. It should be noted that these theories are not necessarily wrong, or even mutually exclusive; given the complexity of building a successful AGI, multiple blockers are bound to crop up along the way. A few of the more prominent theories are

discussed below.

**AGI is impossible:** The idea that humans are inherently different from other animals, and possess a "soul" or "essence" distinct from ordinary matter, has a very long history, going all the way back to ancient Egypt. Although the idea of an immaterial soul rarely appears in modern science, the idea that no computer program can have a mind in the same sense that humans do is alive and well (Searle 1980). However, for practical purposes, we need not debate the philosophy of whether computers can truly think, which is still being discussed more than fifty years after the question was originally raised by Turing and his colleagues (Turing 1950). The most important matter under consideration is the influence that a successful AGI will have on the world, and it is very widely agreed that AGI can, in principle, have the same kinds of effects as a thinking human; after all, it is always theoretically possible to just simulate the entire brain atom-by-atom.

**AGI is computationally impossible:** Penrose and Hameroff, among others, have claimed that, since all thought requires interactions in quantum physics present in the human brain (but not in classical computers), no amount of classical computation will ever be able to replicate the important features of the human brain (Penrose 1991). Although this is certainly possible, it does not make AGI physically impossible or even beyond human capability; we just have to use quantum hardware for AGI projects, instead of standard Turing-equivalent classical hardware. This approach has not been investigated in detail due to the structural problems with AGI research, as scientists are unlikely to be able to publish as many papers on something as complex as quantum physics as colleagues can publish on more conventional ideas. Hence, even if quantum computing is the only viable pathway to AGI, this does not explain why AGI has failed to make progress, and the incentive structure in the field is still a necessary and sufficient explanation.

**Lack of computing power:** Transhumanist advocates, such as Ray Kurzweil and Hans Moravec, often cite a lack of computing power as one of the main reasons why AGI and, more generally, AI, have failed to take off over the past several decades (Moravec 1998). Specific kinds of narrow AI, such as chess-playing programs, speech recognition, and OCR, have indeed shown improved performance over the past two decades, as the cost of computing power has dropped. But all of these sub-disciplines of AI rely on a large body of accumulated, tested, agreed-upon knowledge, in addition to simply throwing additional computing power at the problem. Even in integer factorization, a field where access to large amounts of computing power is obviously of central importance, algorithmic progress over the past thirty years has increased our capabilities more than progress in hardware (Yudkowsky 2007). As of this writing, no comparable central body of knowledge exists in AGI, as the incentive structure is such that it is in few people's interests to create or contribute to such a thing. It is, of

course, possible to simulate the brain by throwing computing power at software which has largely already been written (Frye 2007), but without a detailed understanding of how the brain's neurons are a specific implementation of a generalized theory of intelligence, modifying such a simulation in any interesting way would be extremely difficult, not to mention dangerous to the person in question.

**Lack of funding:** Lack of funding as a blocker to progress is, essentially, a special case of the misaligned-incentives theory. Relative to government and corporate funding for other forms of basic research- the budget of the National Science Foundation in 2008 was only \$6 billion (National Science Foundation 2008)- artificial intelligence research as a whole is well-funded. IBM's Deep Blue project alone cost over \$100 million, spread over a span of twelve years (Herardian 2006). However, because the vast majority of the incentives that determine funding priorities are set up to reward work on building narrow AI and publishing additional papers, it is still very unlikely that a serious, unified, long-term project aimed at building AGI will be funded. Such projects have, of course, been tried numerous times in the past (Feigenbaum 1983), but because the big rewards in AGI are almost always dependent on achieving end goals and not on intermediate progress, funding is usually cut off when the end goals are not achieved on schedule even when significant progress is made.

## Achieving AGI

It is, of course, true that building an AGI is an extremely difficult problem; the human brain is the most complex object yet studied by science (Shepherd 1994), and we are trying to replicate its abilities on hardware which is radically different in character from the networks of nerve cells employed by biological brains. To separate out the problems stemming from the complexity of the subject matter from the problems caused by misaligned incentives, it should be helpful to look at the history of the neighboring field of neuroscience. Like artificial intelligence, neuroscience has produced a huge number of theories and ideas over the past fifty years; like AI, neuroscience is nowhere near complete, and still has a great deal of complexity to unravel. However, it seems that AGI has yet to produce any core body of knowledge which is agreed upon by everyone- there is no single theory of how to build strong AI, or a good-sized piece of strong AI, which has long since been accepted by researchers and is now firmly established enough to teach to college students in textbooks. Neuroscience, meanwhile, has produced such theories in spades, and the body of knowledge which has been slowly accumulating (Wilson 2001) is far too vast to be discussed, or even skimmed over, in any one paper.

The primary difference between AGI and neuroscience, physics, biology, or other well-established fields, is, then, that it has yet to establish a set of basic, agreed-upon theories which can then be built off of and expanded. Like

AGI, the incentive structure of physics, chemistry, etc., is largely focused on producing short-term results and publishing lots of papers. However, because each of these fields has a well-established knowledge base, each new, good idea is accepted into the ever-growing general body of knowledge, rather than being discarded or forgotten. Even though the reward system still favors short-term goals, the longer-term goals also get accomplished as a byproduct of effort for short-term goals; for example, the goal of alchemy was finally achieved in the early 20<sup>th</sup> century, but only after a hundred years of building up a large, well-understood common body of knowledge and a sufficiently large engineering base.

There are, therefore, two primary ways of producing a working AGI: either realign researcher incentives on a large scale to favor work on long-term goals over work on short-term goals, or build up a basic, unified theory of general intelligence, which can then be added to piecemeal by individual researchers and small groups. Neither one, of course, will be easy, but once either one is accomplished, there will finally be sustainable momentum in the direction of solving the AGI problem.

## The Manhattan Project Approach

At the present time, the primary obstacle to launching a large-scale, long-term AGI effort, in the style of the Manhattan Project or the Human Genome Project, is the lack of sufficient funding. It is critical that the sponsors of such a project maintain a sustained emphasis on the long-term goal of AGI, instead of short-term or intermediate goals; numerous large-scale AGI projects have been funded in the past, only to have the plug pulled when the financial backers realized the difficulty of AGI and started to see short-term goals as more attractive. Funding for research, in the United States and most Western countries, is provided primarily by government institutions and private, for-profit corporations. Reallocating funding to serious, long-term AGI projects will therefore require a fairly drastic change in the types of activities funded by these entities, which will, in turn, require a change in their internal motivations. Several possible routes are discussed below, but there are undoubtedly more which have not yet been looked at.

**The Lobbyist Approach:** Directly lobbying the leaders of government agencies, corporations, not-for-profit groups, and other entities is a tried-and-true approach, but in the case of AGI research, there are numerous obstacles to be overcome. In the United States alone, there are already more than thirty thousand registered lobbyists, with a collective budget of over \$2 billion (Bimbaum 2005), all of whom compete for attention from the same, extremely small pool of legislators and other officials. All large corporations, universities, and not-for-profit groups also have numerous special interests vying for their support. The primary motivator for AGI research- the potential for huge, civilization-altering benefits if strong AGI is achieved- has already been overplayed by the previous generation of researchers, and the long history of failed

projects has previously lead both DARPA and European governments to cut off funding to AI more than thirty years ago (Hendler 2008). Even computer scientists and programmers, to say nothing of politicians and bureaucrats, are now leery of being associated with AI (Markoff 2005).

**The Competition Approach:** The Space Race, from 1957-1973, generated few immediate benefits for either the United States or the Soviet Union. However, because it was seen as a matter of patriotism and national pride, both governments poured money into their space programs, eventually spending over \$100 billion in inflation-adjusted dollars (Stine 2008). This came about, however, in large part because progress in the Space Race was seen as a symbol of national engineering and scientific prowess. A successful AGI project will almost certainly be as difficult, or even more difficult, than landing men on the Moon, but it's more than likely that there will be no obvious milestones on the way. Unlike the Space Race, where new missions and new accomplishments were announced on a regular basis, it seems to be entirely possible that an intelligence which is 99% complete will just sit there and not do anything interesting. It would be difficult, if not impossible, to get the broader public to view a project that doesn't have any obvious intermediate progress markers as a key measure of national technological aptitude.

**The War Approach:** War does has a long history as a strong motivator for research; for example, although the potential for a nuclear fission chain reaction was discovered by Szilard in 1933 (Szilard 1936), it was not until 1939, after Germany had started conducting research into the feasibility of a nuclear weapon, that Einstein's letter to Roosevelt convinced him to support government research into nuclear reactions (Einstein 1939). However, even though DARPA and other military agencies have provided funding for AI in the past, and continue to do so, their efforts are, by the necessity of military politics, focused on short-term, militarily-relevant applications and narrow AI (DARPA 2008). Once again, because of the high likelihood of intermediate stages with no obviously beneficial effects, it will be difficult, if not impossible, to portray AGI as a war-winning weapon.

### **The Lavoisier Approach**

Over the past 2500 years, the motivations behind research and the sources of research funding have changed dramatically, but the psychology of the humans undertaking the work of discovery has remained essentially constant (Brown 2000). Large-scale, Manhattan Project-style efforts are historically quite recent, and given that there was no overarching central hierarchy, we should expect the same psychological factors present today to lead to similar problems in historical times. Nevertheless, it is clear from the vast amounts of progress made in science that many fields seem to have escaped stagnation even though the psychology of the researchers involved remained mostly the same. What finally enables individual fields to make large amounts of progress, even though they

aren't focused on a central goal, seems to be the discovery of a set of underlying, easily understood, universally accepted principles, which can then be expanded upon and modified as necessary. Historical examples abound; in astronomy, the discovery of gravity and the laws of orbital dynamics in the 17<sup>th</sup> century; in chemistry, the discovery of conservation of mass and the modern theory of elements in the late 18<sup>th</sup> century; in genetics, the unraveling of the DNA -> RNA -> protein -> trait mechanism, called the 'central dogma of molecular biology' (Crick 1970), during the 1950s and 60s. More recently, a similar shift seems to have happened in the social sciences, with the introduction of evolutionary psychology and the field of heuristics and biases (Barkow 1996).

However, unlike many of these other fields, intelligence involves a large number of dissimilar, interacting components, and it seems very unlikely that the underlying mathematics will be as simple as conservation of energy, or even as simple as quantum field theory (Yudkowsky 2006). While many of the basic principles of chemistry and physics were discovered by Lavoisier and Newton, respectively, the complexity of the subject matter implies that creating a basic framework for AGI will require work by a relatively large number of scientists, in order to bring together the large number of missing puzzle pieces. This bodes ill for the future of the field, as scientists with both the capability to contribute to fundamental theories and the desire to do so seem to be quite rare; for example, Kepler only discovered the laws of celestial mechanics in the 17<sup>th</sup> century, even though the motions of planets had been described in detail by Ptolemy over a millennium earlier, and his model had been in wide use among scholars since that time. Nevertheless, this approach has the advantage of not requiring large amounts of funding or the support of any large institution.

### **Avoiding Stagnation**

After an AGI community is solidly established, the most important long-term threat to progress seems to be from bureaucratic sclerosis. Although sclerosis is, like the problems currently present in the field, caused by a failure of the researchers' incentive structure, large bureaucracies will wind up causing different kinds of failures than funding agencies. Instead of producing lots of poorly-integrated new ideas, big organizations tend to favor internal marketing and political maneuvering- over half the time of the average employee of a large organization is dedicated to such activities (Whitney 2004). Under some circumstances, a large AGI community might wind up producing even less new research than the relatively small community of the past fifty years, if there is a large enough drop in per-capita productivity from all the wasted time.

The primary driver of stagnation in large organizations seems to be that the size and complexity of the organization isolates individual employees from being worried about overall progress towards the organization's goal. This causes them to focus their efforts towards

improving their own, internal rank in the hierarchy, and eventually results in the entire organization losing its original focus, unless a strong external pressure is applied. From the historical record, it is fairly clear that the time required for sclerosis is usually on the order of several decades; in the shorter term, the Manhattan Project largely avoided it, and NASA avoided it during the 1960s (Zubrin 2003). But because the former was dissolved and the latter was not, NASA eventually did develop a stifling bureaucracy, in spite of their technological aptitude and their efforts to prevent it (Feynman 2001). For purposes of brevity, methods of preventing stagnation which have already been tried and found not to work on a large scale—such as NASA's practice of hiring former engineers as managers— are omitted in the discussion of possible solutions below.

**Replacing all workers:** The obvious solution to the problem of stagnation is to split up the organizations working on AGI after a few years, but needless to say, this is not a practical way of doing business; either the new organizations will carry over most of the structure from the old organizations through collective inertia, or a great deal of project-specific talent will be lost in the mess. In theory, it would be possible to simply fire people or assign them to new projects after a few years, a proposal that has been previously suggested for avoiding sclerosis in software companies (Papadimoulis 2008). But while software developers are largely seen as interchangeable, professional researchers usually become highly specialized, and there is a great deal of evidence that they do their best work only when working with a very select subset of their peers (Root-Bernstein 1997), making this approach counterproductive.

**Replacing unproductive workers:** The tenure system at most universities guarantees that the vast majority of senior professors will always have a position, regardless of their performance or what they do research on, which causes older ideas and methods to be given excessive weight. An alternative option is to simply not grant tenure, and fire scientists along with administrative staff if they can't carry their weight. Although this would work well in principle, in practice, it runs into the same problem that led to the creation of the tenure system in the first place—performance is usually judged by managers and other scientists in the field, and so if a single idea or set of ideas comes to dominate the memospace, those who disagree with it will find it hard to get funding or peer recognition. Indeed, in spite of the tenure system, there is some evidence that this is already happening in the theoretical physics community with string theory (Smolin 2007, Woit 2007).

**Removing management:** Since most of the difficulties with bureaucratic sclerosis originate from managers and administrators, not scientists, it seems fairly obvious that if administrators were never hired in the first place, such difficulties would be markedly less pronounced. In 1970, this would have been unthinkable; scientists generally need to be among their peers in order to do top-notch work

(Root-Bernstein 1997), and since this implied physical proximity, there had to be some kind of administrative structure to manage the resulting organization. However, with the advent of the Internet, it seems at least possible that researchers with sufficient initiative could bypass the traditional academic system and avoid administrators entirely. This already seems to be the trend in many fields of science, with more and more papers being published in online archives instead of traditional scientific journals (Jackson 2002).

## Hope for the Future

If the current AGI community succeeds in creating a dynamic, innovative, productive field surrounding the idea of strong AI, the human species will reap the benefits of true artificial intelligence far sooner than we would have otherwise. But because most research shares roughly the same organizational structure and the same psychological motivation, we should be able to take the lessons of artificial intelligence and apply them to other fields which may currently be languishing in obscurity. Although this will be difficult, we have no reason to think that it is impossible— we have only developed the communication networks required to coordinate scientific revolutions relatively recently, and so unlike in the field of AGI proper, there is no history of failed attempts to discourage us.

## References

- Barkow, Jerome H., Leda Cosmides, and John Tooby, eds. The Adapted Mind: Evolutionary Psychology and the Generation of Culture. New York: Oxford UP, 1996.
- Bimbaum, Jeffrey H. "The Road to Riches Is Called K Street." Washington Post 22 June 2005: A01.
- Brown, Donald E. "Human universals and their implications." Being Humans: Anthropological Universality and Particularity in Transdisciplinary Perspectives. By Neil Roughley. Danbury: Walter De Gruyter Incorporated, 2000.
- Crick, Francis. "Central dogma of molecular biology." Nature 227 (1970): 561-63.
- Einstein, Albert. Letter to Franklin D. Roosevelt. 2 Aug. 1939. Peconic, Long Island, NY.
- Feigenbaum, Edward A., and Pamela McCorduck. The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World. New York: Addison-Wesley Longman, Incorporated, 1983.
- Feynman, Richard Phillips, and Ralph Leighton. What Do You Care What Other People Think?: Further Adventures of a Curious Character. Boston: W. W. Norton & Company, Incorporated, 2001.
- Frye, James, Rajagopal Ananthanarayanan, and Dharmendra S. Modha. Towards real-time, mouse-scale cortical simulations. Rep.No. RJ10404. Research Division, IBM.

- Green, Leonard, Joel Myerson, and Pawel Ostaszewski. "Discounting of delayed rewards across the life span: age differences in individual discounting functions." Behavioral Processes 46 (1999): 89-96.
- Hendler, James. "Avoiding Another AI Winter." IEEE Intelligent Systems Mar.-Apr. 2008: 2-4.
- Herardian, Ron. "Project Deep Blitz: Chess PC Takes on Deep Blue." ExtremeTech. 26 Jan. 2006. Ziff Davis Publishing Holdings. 21 Oct. 2008 <<http://www.extremetech.com/article2/0,1697,1915575,00.asp>>.
- Hofstadter, Douglas R. Metamagical Themas: Questing for the Essence of Mind and Pattern. New York: Basic Books, 1996.
- Howe, Jim. "Artificial Intelligence at Edinburgh University: a Perspective." School of Informatics. June 2007. Edinburgh University. 21 Oct. 2008 <<http://www.inf.ed.ac.uk/about/aihistory.html>>.
- Jackson, Allyn. "From Preprints to E-prints: The Rise of Electronic Preprint Servers in Mathematics." Notices of the American Mathematical Society 49 (2002): 23-32.
- Jaynes, E. T. Probability Theory: The Logic of Science. Ed. G. Larry Bretthorst. New York: Cambridge UP, 2003.
- Kurzweil, Ray. The Singularity Is Near: When Humans Transcend Biology. New York: Viking Adult, 2005.
- Landsburg, Steven E. The Armchair Economist: Economics and Everyday Experience. New York: Simon & Schuster, Incorporated, 1995.
- Latour, Bruno, and Steve Woolgar. Laboratory Life: Construction of Scientific Facts. Ed. Jonas Salk. New York: Princeton UP, 1992.
- Lighthill, James. "Artificial Intelligence: A General Survey." Artificial intelligence: A paper symposium. London, UK: Science Research Council, 1973.
- Markoff, John. "Behind Artificial Intelligence, a Squadron of Bright Real People." The New York Times 14 Oct. 2005.
- Moravec, Hans. "When will computer hardware match the human brain?" Journal of Evolution and Technology 1 (1998). Journal of Evolution and Technology. Dec. 1997. Institute for Ethics and Emerging Technologies. 21 Oct. 2008 <<http://www.transhumanist.com/volume1/moravec.htm>>.
- National Science Foundation. Office of Legislative and Public Affairs. "President Signs Omnibus Appropriation Bill." Press release. 8 Jan. 2008. 21 Oct. 2008 <[http://www.nsf.gov/about/congress/110/highlights/cu08\\_0108.jsp](http://www.nsf.gov/about/congress/110/highlights/cu08_0108.jsp)>.
- "Overview." Urban Challenge. DARPA. 21 Oct. 2008 <<http://www.darpa.mil/grandchallenge/overview.asp>>.
- Papadimoulis, Alex. "Up or Out: Solving the IT Turnover Crisis." The Daily WTF. 29 Apr. 2008. 22 Oct. 2008 <<http://thedailywtf.com/articles/up-or-out-solving-the-it-turnover-crisis.aspx>>
- Penrose, Roger. The Emperor's New Mind: Concerning Computers, Minds, and the Laws of Physics. New York: Penguin (Non-Classics), 1991.
- Phillips, Eve M. "A Commercial Look at Artificial Intelligence Startups." Thesis. 7 May 1999. 21 Oct. 2008 <<http://www.sts.tu-harburg.de/~r.f.moeller/symbolics-info/ai-business.pdf>>.
- Root-Bernstein, Robert Scott. Discovering : Inventing and Solving Problems at the Frontiers of Scientific Knowledge. New York: Replica Books, 1997.
- Russell, Stuart J., Peter Norvig, and John F. Canny. Artificial Intelligence: A Modern Approach. Upper Saddle River: Pearson plc, 1999.
- Searle, John R. "Minds, Brains, and Programs." Behavioral and Brain Sciences 3 (1980): 417-57.
- Shepherd, Gordon M. Neurobiology. New York: Oxford UP, 1994.
- Smolin, Lee. The Trouble with Physics : The Rise of String Theory, the Fall of a Science, and What Comes Next. New York: Mariner Books, 2007.
- Stine, Deborah D. The Manhattan Project, The Apollo Program, and Federal Energy Technology R&D Programs: A Comparative Analysis. Rep.No. RL34645. Resources, Science, and Industry Division, Congressional Research Service.
- Szilard, Leo. Improvements in or relating to the transmutation of chemical elements. British Admiralty, assignee. Patent GB 630726. 1936.
- Testimony of Dr. Robert Zubrin to the Senate Commerce Committee, Oct 29, 2003*, 108th Cong. (2003) (testimony of Robert Zubrin).
- Turing, Alan M. "Computing Machinery and Intelligence." Mind 59 (1950): 433-60.
- Whitney, John O. The Trust Factor: Liberating Profits and Restoring Corporate Vitality. New York: McGraw-Hill Companies, The, 1994.
- Wilson, Robert, and Frank Keil, eds. The MIT Encyclopedia of the Cognitive Sciences (MITECS). New York: MIT P, 2001.
- Woit, Peter. Not Even Wrong: The Failure of String Theory and the Search for Unity in Physical Law. New York: Basic Books, 2007.
- Yudkowsky, Eliezer. "Artificial Intelligence as a Positive and Negative Factor in Global Risk." Global Catastrophic Risks. By Martin J. Rees. Ed. Nick Bostrom and Milan Cirkovic. New York: Oxford UP, Incorporated, 2008.
- Yudkowsky, Eliezer. "Introducing the Singularity: Three Major Schools of Thought." Singularity Summit 2007. Palace of Fine Arts Theater, San Francisco. 8 Sept. 2007. Future Current. 24 Oct. 2007. Accelerating Future. 21 Oct. 2008 <<http://www.acceleratingfuture.com/people-blog/?p=209#more-209>>.
- Yudkowsky, Eliezer S. "Levels of Organization in General Intelligence." Artificial General Intelligence. Ed. Ben Goertzel and Cassio Pennachin. New York: Springer, 2006.
- Zeev, Alfassi. Stat Treatm of Analy Data. New York: Taylor & Francis Group, 2004.