



# Modern Research Studies

**Editor-in-Chief**  
**Gyanabati Khuraijam**

**An International  
Journal of  
Humanities and Social  
Sciences**

*An Indexed & Refereed e-Journal*

[www.modernresearch.in](http://www.modernresearch.in)

**Title: The Possibility of Creating a Conscious Entity:  
A Humanities Perspective**

**Author/s: Rob Harle**

**Volume 1, Issue 2  
September 2014**

**pp. 112–128.**

Disclaimer: The views expressed in the articles/contributions published in the journal are solely the author's. They do not represent the views of the Editors.

## **The Possibility of Creating a Conscious Entity: A Humanities Perspective**

**ROB HARLE**

Independent Scholar

29 Basil Road

Nimbin NSW 2480

Australia

**Abstract:** The purpose of this essay is to determine, in principle, if it is possible to create a conscious machine. That is, an artificial intelligence within a non-carbon based entity which has the qualities of intentionality, conscious awareness, learning ability and the ability to develop intellectually in a similar way to that of a human being through socialisation and enculturation. The paper discusses the notion of human intelligence and the importance of socialization to develop this intelligence using the concept of Artificial Intelligence as a reference model. The investigation throws light on the nature of the human unconscious mind and concludes that in principle the creation of an artificially intelligent entity is possible.

**Keywords:** artificial intelligence, enculturation, human consciousness, embodiment, social interaction, intentionality.

The purpose of this essay is an attempt to determine, in principle, if it is possible to create a conscious machine. I would neither discuss the moral or ethical issues such an entity would raise, nor the desirability of such an enterprise. I also do not want to discuss the minute details of AI (Artificial Intelligence) methods – programming, computer architecture and so on. Apart from the fact that I’m not qualified to do so, getting bogged down in detail can often obscure or confound the overall picture. Also, this essay is more of an exploration of the issues involved and a general introduction to the rather esoteric field of AI, in which it seems there are as yet no definitive authorities.

The notion of creating an artificial entity, that is, ‘a human created in the image of a human’ has been around a lot longer than the last thirty or forty years of the Computer Age. Mary Shelley’s *Frankenstein* dealt with the ‘scientific’ creation of a human-like entity on many levels. One of these levels was socialisation and in an indirect, somewhat ironic way is relevant to the possible creation in the early twenty-first century of an intelligent, conscious silicone-digital-molecular ‘creature’. I discuss the socialisation aspect of AI in detail furtheron.

The creation of computational devices and those with a ‘memory’ goes back many thousands of years. The Chinese invented the Abacus some 5,000 years ago, the water clock around 3,000 years. The first truly mechanical clock, which ticked, was built in 725CE. In 1642, Pascal invented the world’s first automatic calculating machine, the Pascaline. Shortly after in 1694, a computer was invented by Leibnez.

Weaving was to see many mechanical and computational innovations at the beginning and during the Industrial Revolution. The Jacquard loom, around 1805, was operated by punch cards. This was a significant ‘advance’ in transferring instructions from a human to a machine to produce a product, in absentia. This was of course the precursor to punch card computers of the mid-1900s.

Charles Babbage developed the Analytic Engine, the world’s first ‘real’ computer. Ada Augusta published her own notes on this engine and ‘speculated on the ability of computers to emulate human intelligence’ (Augusta 1842). The year was 1842. Hollerith perfected the automatic punch card tabulating machine and founded a company which later became IBM.

In 1921, the word “robot” was coined by Czech dramatist, Capek. His science fiction drama *Rossum’s Universal Robots* describes how intelligent machines, originally servants of humanity, end up taking over and destroying their creators.

In 1937, Turing developed his famous Turing Machine, a theoretical model of a computer. In 1950, he developed his almost infamous, “Turing Test” to assess the intelligence of a machine compared with that of a human. The usefulness of this test is now seriously doubted (see Searle and Hofstadter). From about this time on computers simply became faster and with greater and greater capacity and problem solving capabilities (Kurzweil 1999, 261-280).

Relays gave way to vacuum tubes; these gave way to transistors, which in turn gave way to the silicon chip. This ubiquitous ‘chip’ may soon give way to a new technology known as molecular organic circuitry. Silicon chips are reaching maximum efficiency due to physical limitations, a custom-designed molecule called rotaxane, may replace the old style digital silicon computer with a molecular computer within five years (Chang 1999). This very brief history helps ground the following discussion and shows the progressive nature of human invention. In a sense it helps justify the predictions of futurists, such as Kurzweil, that machines will be conscious and equally as intelligent as humans within thirty years (Kurzweil 1999).

Before I address the question of artificial consciousness, we need to have a clear idea of what constitutes ‘natural’ human consciousness. This is no easy task, which is evidenced by the monumental amount of literature devoted to the “race for consciousness”. Searle defines consciousness as, “...a biological feature of human and certain animal brains. It is caused by neurobiological processes and is as much a part of the natural biological order as any other biological features such as photosynthesis, digestion, or mitosis” (Searle 1992, 90).

Hobson believes, “...the mind is all the information in the brain” and defines consciousness as, “...the brain’s awareness of some of that information” (Hobson 1994, 202-204). Hobson also believes that the brain-mind is a unified system; they are inextricably linked (Hobson 1994, 26). Hobson’s theory is supported by a large amount of testable, neuro-physiological, experimental evidence.

Drawing on this evidence and the ideas of Searle, Dennett and Gelernter, I believe when we refer to ‘the mind’ we are referring to a highly complex system, which is a combination of electrochemical-neural interactions in the organ within the skull called the brain. The brain interacts with the physical body, and through sensory input/output, the environment external to it. Embodiment with sensory input is an *essential* requirement for a mind to exist. This sensory input further allows for another essential requirement for the creation and maintenance of mind, and that is socialisation.

In a normal human adult, the unified brain-mind system has both non-conscious mental states (memories, regulation of respiration, etc.) and at times conscious mental states. Consciousness is simply one state of this unified brain-mind system. Consciousness is largely controlled by the brain’s chemical system known as the aminergic-cholinergic system.

The aminergic system (amines) governs our waking state and the cholinergic (acetylcholine) system governs our dreaming state. These systems are in dynamic equilibrium and neither one is ever totally inactive. The ratio of these chemicals can now account for many previously mysterious states of consciousness, such as hypnosis, dementia and fantasy. As we approach sleep, the cholinergic chemical increases and maintains dominance whilst asleep. As we wake up normally, the reverse happens and the aminergic system becomes dominant. If we are awoken suddenly, we temporarily experience confusion and disorientation because the chemical system needs a little time to re-establish its correct ratio/balance for the respective, consciously desired states (Hobson 1994, 14-16). I have noted a most interesting correlation in the work of Gelernter. He believes mental focus moves from high to low. At the high focus end, we are most alert, logical and deal with step-by-step problem solving. At the low focus end, that is, as we move down the spectrum we do not think logically, our minds move easily from one unrelated subject to another. Creative solutions to problems occur at this level, the ones that have previously defied logical solution. It is at this level that inspiration suddenly hits us. Further down the spectrum the onset of sleep and then dreaming occurs. We must bear in mind that during REM sleep, we dream; the

awareness of dreams or dream fragments, even though we are asleep, is still part of a conscious state. This description of mental states fits in perfectly with the action of the aminergic-cholinergic system.

I can see that some philosophers, though happy enough with the above, may still argue that it does not say what consciousness actually is. Searle helps overcome this conundrum through his belief that our vocabulary and consequently our mode of thinking is at fault. It is incorrect to think that a state must be either mental or physical. Searle believes such apparent oppositions as these are false as, “Consciousness is a mental, and therefore physical property of the brain in the sense in which liquidity is a property of systems of molecules, eg.  $H_2O$  (Searle 1992, 14). Further, “...consciousness *qua* consciousness, *qua* mental, *qua* subjective, *qua* qualitative is physical, and physical because mental (Searle 1992,15). This approach I believe is plausible in answering the elusive question of what consciousness *is*.

Dennett, discussing Jaynes and Nagel, describes the chasm between inert matter and the inwardness of a conscious being, in the example of brick and bricklayer (Dennett 1998, 122). If we accept that a brick cannot be conscious, and it is by no means a universally accepted conclusion, many tribal societies believe inanimate objects such as a stone do have a kind of consciousness or at least spiritual essence. In principle, we must remain open to this because we are not stones so can never ‘really’ know what it is like to be a stone. If we do not accept this possibility and insist that a stone is inert and a stonemason is conscious how can this be? Suppose we ‘deconstruct’ both stone and stonemason. Prior to total deconstruction we get down to molecules and atoms, say carbon, silica, hydrogen and so on – the very same building blocks fundamentally present in both stone and stonemason. Further on, we arrive at Quantum states, probabilities, particles and waves. So, how and where from does this ‘oneness’ consciousness become an attribute of the stonemason and not the stone? I believe that Hobson and Searle are correct in insisting that consciousness arises from brain-mind states. The reason it can arise is because a *functioning system*, with just the right attributes, causes it to exist. It seems clear from scientific ‘deconstruction’ of stones that they do not have a brain-mind system and therefore cannot be conscious in any sense that humans are.

Franklin argues that mind is graded not Boolean; this fuzzification of mind *allows for* some degree of mind in animals and possibly machines, though it may be in the mechanical sense, that is, without qualia (Franklin 1995, 412). Perhaps the degree of consciousness is proportional to the complexity of the system from which it arises. Hence we might imagine a consciousness complexity scale from one to one hundred, along which a plant may be zero, an ant two, a dog sixty, an ape seventy and a human ninety. These figures are of course speculative but it helps illustrate the point. This consciousness scale has nothing to do with Gelernter's low-high focus of consciousness. His model applies separately to each species which is conscious.

Rather than providing a precise definition of consciousness, in the foregoing pages, I have attempted to approach the phenomena from various angles, to at least find some things consciousness is not and others that must be present for consciousness to arise. Two characteristics of consciousness that are particularly relevant to this discussion are awareness and intentionality. Without awareness we are not conscious, more precisely we are not conscious 'of' something. Awareness may be of external events or internal brain-mind mentation, such as the dreams of REM sleep. Various meditational states, like dreams, have absolute minimal external stimulus yet the individual may be conscious of these, as an example, consciousness of 'nothingness'. Austin explains many of these mysterious mental phenomena in a large body of research work, represented in *Zen and the Brain*.

For an agent to be considered conscious it must display Intentionality. As Searle points out, this does not mean that intentionality is consciousness though. "Intentionality is that property of many mental states and events by which they are directed at or about or of objects and states of affairs in the world" (Searle 1983, 1). Intentionality needs to be divided into intrinsic and 'as-if' intentionality for clarification and for consideration in an artificially intelligent entity.

If a person makes a statement such as, "I am afraid of snakes", this is an example of intrinsic intentionality. If your personal computer

displays a message, ‘I am afraid of snakes’ this is, as-if intentionality. Many devices display as-if intentionality, a thermostat is a good example, but none of these devices ‘so far’ have the presence of any mental phenomena. If an office thermostat controlling the air conditioning had a speaker attached and always report when it was turning up the heating and then one day said, “I’m not turning up the heating today because I’ve had it with you people treating me like a dumb wallflower!” Then it would be displaying intrinsic intentionality. Deciding between as-if intrinsic intentionality and intrinsic intentionality will be a criteria for assessing true artificial intelligence and will not be as easy as it may seem.

For an entity to be considered consciously aware that it exists, it must possess intrinsic intentionality. If consciousness and intentionality are features of a unified brain-mind system, we need to ask how brain-mind states come about, over and above their chemically controlled basis. If the brain-mind is simply an ‘information-processing’ organ, with on-off switching and access to a huge knowledge database (memory), that is, computational power plus knowledge, we would already have developed rudimentary artificial intelligent conscious machines. It seems, in principle, this approach to AI is fundamentally flawed.

One reason for this is that all the information in the world is not the measure of intelligence; one measure of intelligence is the ability of an organism to function within its environment and survive the normal hazards of that environment.

A walking encyclopedia will walk over a cliff, for all its knowledge of cliffs and the effects of gravity, unless it is designed in such a fashion that it can find the right bits of knowledge at the right time, so it can plan its engagement with the real world. (Dennett 1998, 194-195)

The failure of traditional or classical AI led to the development of the connectionist paradigm, this included neural networks operating in parallel, similar to the way the brain operate and accesses subsystems



which ‘do their own thing’ at a local level. The connectionist model allows a system to learn and expand its program as it encounters various situations. Whereas classic AI (rule-and-symbol-based) is good at logic and long term planning, it is inadequate for real-time motor control and perceptual recognition (Clark 1997, 59).

Classic AI often sees the computer as analogous to the human brain; little wonder scholars such as Black argue that, “The digital computer analogy is fatally misleading” (Black 1991, 3). Black further argues that the software-hardware dichotomy is artificial, “...software and hardware are one and the same thing in the nervous system” (3). Whilst Black is no doubt correct in that DNA instructions are encoded right in the cell, ‘on site’ ready to do their job, I do not agree they are “one and the same thing”. They are instructions embedded into the molecular matrix of particular parts of the cell.

Software instructions are part of a machine’s memory, embedded electrically in the matrix of the memory medium. In all but the oldest AI programs, these instructions form part of feedback loops which modify and expand the original instructions and also rearrange their positions and relevance in the software’s hierarchical structure. Further to this, it has always struck me as naive and almost absurd that AI researchers, up until a few years ago, imagined they could create an artificially intelligent machine and do so with the machine in isolation, simply by increasing speed and adding more computational power. This procedure has resulted in very powerful machines which can outperform humans in many respects; however, this has nothing to do with intelligence. The nurturing period of a human infant, with the longest neoteny of any species, together with the interaction of infant with other infants and adults is partly the basis of human intelligence. During this period, up to three years of age, the development of the limbs, the structuring of neural pathways and the gradual appreciation by the infant that it is an autonomous agent all take place. In my opinion, without an equivalent period of infancy, no machine will ever even approximate human consciousness.

I think much wasted discussion and programming effort has taken place because of the limited vision of just what an AI entity

would require prior to being able to develop internal states that could give rise to consciousness. The coming into existence of the World Wide Web may be a great benefit for the socialization of AI entities; entities could be in-line for extended periods and use the Web as a classroom for learning facts and as a place for social interaction. I cannot speculate on the benefits of Virtual socialisation over real socialisation but there are futurists who regard it as equally efficacious. Vinge, in a paper presented to NASA, discussed the possibility of computer networks ‘waking up’; the Web with its millions of users at any one time could possibly be regarded as an intelligent, ‘artificial’ entity in its own right (Vinge 1993).

Regardless of how an entity experiences socialisation, its first requirement is the embodiment, that is, the ‘mind’ part of the entity must have some sort of physical attributes which help locate it spatially. An entity cannot be aware of its existence unless it has reference to other objects which are *not it*. Thelen and Smith have done some important pioneering work in developmental psychology which has been especially relevant in dispelling the entrenched Cartesian notion of ‘mind’ as a separate, controlling homunculus like thing. Known as the Dynamic Systems approach to the development of cognition and action, this approach has proven beyond all doubt that various parts of a system ‘do their own thing’. Literally, the brain does not know how to do certain things nor that they have occurred, such as some of the aspects of an infant learning to walk (Thelen and Smith 1994, Chap.1). These discoveries have major implications for AI, the low-level design of the body of an entity allows for local knowledge and control without the burden of complex, resource hungry, central executive control.

This fundamental approach has been implemented in MIT’s Cog Project. I was excited to come across this work as it is actually doing what I thought were the minimum necessary basic steps in creating an intelligent artificial entity. As well as utilising local feedback control, the process of socialisation is being carried out equally with hardware (body) modification and programming evolution.

Before describing this project in detail it is worth noting that Collins, towards the end of the eighties, was insisting that socialisation and enculturation are essential components of intelligence. Although he did not describe the way machines must be socialized to exhibit intelligence, his work pre-empted what is happening with the CogProject. Collins made a very interesting point which few people seem to think about and that is, perhaps, rather than humans trying to make machines like humans, they are becoming more like machines themselves. Quoting Dreyfus, he says that “Our risk is not the advent of super-intelligent computers but of sub-intelligent human beings” (Collins 1990, 190).

The Cog Project at MIT, under the direction of Rodney Brooks, is an ongoing concern which seeks to build “human-like artificially intelligent systems”, not systems which master a single domain but those which can adapt to many complex tasks in the real world in real time. This goal has led to the rejection of many of the procedures in classical AI and also of the assumptions about human intelligence which are features of this discipline.

The guiding principle of the Cog team, is that, “...human intelligence is a direct result of four intertwined attributes: *developmental organization, social interaction, embodiment and physical coupling, and multi-modal integration* (Brooks et al. 1998). Before discussing these attributes in detail I will describe the assumptions about intelligence which classical AI still believes, and Brooks et al, eschew – monolithic internal models, monolithic control and general purpose processing.

*Humans have no full monolithic internal models.* When performing a copying task we do not build an internal model of the complete scene we are attempting to copy. Experiments have shown that, “...humans tend to only represent what is immediately relevant from the environment and those representations do not have full access to one another” (Brooks et al. 1998).

*Humans have no monolithic control.* Evidences from cognitive science, whilst acknowledging control structures, finds no support for a

single unitary control system. Observation of various split brain patients suggests, "... that there are multiple independent control systems, rather than a single monolithic one" (Brooks et al. 1998).

*Humans are not general purpose.* Despite the conventional, commonsense view that humans are equally good at any tasks they attempt, experiments have shown this to be false. The way information is presented affects the ability to solve problems quite significantly.

"Humans, often do not use subroutine-like rules for making decisions" (Brooks et al. 1998). Quite often, emotional rather than rational factors are the major aspects of decision making. The work of Damasio is seminal in this regard (Damasio 1994.)

These three factors alone shift significantly the approach to designing intelligent machines and assessing the intelligence and consciousness levels of other natural animals. Together with the four previously mentioned attributes or "essences of human intelligence" required in an entity it is little wonder classical AI has not achieved its optimistic goals. *Expert Systems*, such as Weizenbaum's famous "ELIZA" program, impressive as they were, really had little to do with true artificial intelligence (Weizenbaum 1976). Although numerous 'hopefuls' believed these systems did display elements of true AI, I do not think Weizenbaum himself ever made such claims.

Returning now to the essential attributes for the development of intelligence (and the possibility) of conscious awareness I will look first at *development*.

(a) *Development*. Development is the framework within which an infant gradually acquires more and more complex skills. "Humans are not born with complete reasoning systems, complete motor systems, or even complete sensory systems." The earlier developmental processes seem to, "...prepare and enable more advanced forms of behavior to develop within the situated context they provide"(Brookes et al. 1998).

(b) *Social Interaction*. "The presence of a caregiver to nurture a child as it grows is essential. This reliance on social contact is so integrated into

our species that it is hard to imagine a completely asocial human” (Brookes et al. 1998). The ABC, some three years ago, televised some secretly obtained footage of children’s institutions in Russia where very young children were abandoned and had absolutely minimal social and physical contact, and certainly none with a caretaker. The children were assessed by Aid workers to be severely emotionally (and from memory) intellectually, undeveloped. Work with autistic children also gives clues as to the importance of social integration. A number of scholars, including Sacks, Synder and Baron-Cohen work and investigate within this field.

(c) *Embodiment*. As Brookes et al. note the most obvious and clearly overlooked aspect of human intelligence is that it is embodied. There is a direct physical coupling between action and perception without the need for intermediary representation. For an embodied system, internal representations can be ultimately grounded in sensory-motor interactions with the world (Lakoff 1987) (Brookes et al. 1998).

One reason for it being ‘overlooked’ I believe is a religious one. For the last two thousand years the dominant influence on Western thinking has been Christianity. Christianity maintains the ‘flesh’ is unclean, necessary for a time to be sure, but ultimately it is the spirit that matters. Following in this tradition Descartes separated the body and mind, as though the body was more or less irrelevant to the mind. I mention this not just as an aside but because of the pervasive influence of spiritual traditions, both East and West, on the psyche of humanity. If the body is essential to the maintenance and formation of mind and consciousness, it raises very serious problems for such doctrines as reincarnation.

The importance and impact of the realisation that mind and body are not separate, that embodiment is essential to the development of anything that can be considered ‘mind’ has not yet been realised by society at large. The corporeal/mental (spirit-soul-mind) dichotomy is so entrenched in our languages and culture that when it is fully realised that there is no central controlling executive, no esoteric special ‘matter’ that constitutes mind, it will be equivalent, in my opinion, to a

Copernican Revolution, the ramifications of which we can only barely imagine at present.

(d) *Integration*. Just as no executive controls our every function, evidence now suggests that no one sensory input (visual, olfactory) is independent of the others. The huge amount of information that comes from the external environment is processed simultaneously and of course gives us our view of the world. “Stimuli from one modality can and do influence the perception of stimuli in another modality” (Brookes et al. 1998). This means any attempt to create artificial intelligence must take this *dependent* phenomenon into consideration.

The Cog team’s methodology recognises the above attributes because they are important aspects of human intelligence and, “...from an engineering perspective, these themes make the problems of building human intelligence easier” (Brookes et al. 1998).

Apart from the fact that embodiment is a necessary criterion of intelligence, by giving their Cog creations, bodies, it allows humans to interact with the robots in a natural way. Further, “...the effects of gravity, friction and natural human interaction are obtained for free, without any computation” (Brookes et al. 1998). One fascinating and, on reflection, essential attribute of the Cog robots, is that of eyes. The team has recognised the vital importance of eye contact between human infants and their caretakers; and later, eye contact with adults. The robots have specially designed complex eyes which allow this interaction and also enable the robot to visually recognise the various people it interacts with each day.

The development of the system is incremental, that is, the earlier learnt behaviors and so on, “...bootstrap the later structures by providing sub-skills and knowledge which can be re-used” (Brookes et al. 1998). Just like a human infant, the system gradually increases its understanding and gradually becomes able to handle more and more complex problem solving tasks. The important thing to realise with this approach is that it is, “...in stark contrast to most machine learning methods, where the robot learns in a usually hostile environment, and

the bias, instead of coming from the robot's interaction with the world, is *included by the designer*" [my emphasis] (Brookes et al. 1998).

The Cog team's approach does not emphasise enculturation as much as I believe is necessary. Socialisation is not quite the same thing as enculturation and whilst, "Social interaction allows humans to exploit other humans for assistance, teaching and knowledge" (Brookes et al. 1998) this does not necessarily imply that culture is being passed on per se. Culture is arguably as important in the development of human intelligence and consciousness as are biological factors, consequently, an intelligent entity needs to learn and be moulded by cultural inputs so as to be able to communicate to others in that culture. Granted, *part* of this transmission of culture takes place during the normal socialisation of humans.

One last aspect of the search for the necessary fundamentals of intelligence and consciousness is the notion of the Unconscious. A colleague's chance remark, asking me just how one would create the Unconscious, even if it was possible to create the equivalent of the conscious mind, led me on an intense investigation of the Unconscious. To my knowledge, this aspect of human mentation has not been discussed in the AI literature. Similar to the central executive controller concept being exposed as a myth so too have I shown that the notion of the Unconscious, particularly in the Freudian sense, is an artificial construct. The Freudian Unconscious with its supposed sexual, libidinal repressions and its expression, symbolically through the *latent* dream content, does not exist. The unified brain-mind has various mental states, most are *nonconscious* at any one time, the brain-mind may contain painful *suppressed* memories but these are nothing to do with the widely accepted and almost unchallenged existence of the Unconscious (Harle 1999). The removal of the Unconscious from consideration in AI research is one further advance towards creating true artificial intelligence.

In conclusion, I have attempted to present in this paper the broad issues involved in the project of creating an artificial intelligent, conscious entity. I believe this is practically and, in principle, impossible if we follow the path of classical AI, that is, computational

power and huge amounts of knowledge (facts). However, if we pursue the approach developed by the Cog team, and once certain hardware constraints are overcome, especially the creation of massive parallel neural network architectures I can see no plausible argument that denies the possibility of creating an intelligent, conscious, non-carbon-based entity.

## References:

- Augusta, Ada. 1842. Trans. "Sketch of Invented by Charles Babbage." By L.F. Menabrea. *Bibliothèque Universelle de Genève*, October, 1842, No. 82. Accessed June 5 2014. <http://www.fourmilab.ch/babbage/sketch.html>
- Austin, J. A. 1998. *Zen and the Brain*. Cambridge, Massachusetts: MIT Press.
- Ballard, D. H. 1997. *An Introduction to Natural Computation*. Cambridge, Massachusetts: Bradford, MIT Press.
- Black, I. B. 1991. *Information in the Brain: A Molecular Perspective*. Cambridge, Massachusetts: Bradford, MIT Press.
- Brooks, Rodney A., Cynthia Breazeal (Ferrell), Robert Irie, Charles C. Kemp, Matthew Marjanovic, Brian Scassellati and Matthew Williamson. "Alternative Essences of Intelligence." AAAI-98. <http://groups.csail.mit.edu/projects/lbr/hrg/1998/group98.pdf>
- Chalmers, D.J. 1996. *The Conscious Mind: In Search of a Fundamental Theory*. New York: Oxford University Press.
- Chang, K. *Many Pentiums on a Grain of Sand*. Accessed June 5 2014. <http://abcnews.go.com/sections/science/DailyNews/nanocomputer990715.html>



- Clark, Andy. 1997. *Being There: Putting Brain, Body and Mind Together Again*. Cambridge, Massachusetts: MIT Press.
- Collins, H.M. 1990. *Artificial Experts: Social Knowledge and Intelligent Machines*. Cambridge, Massachusetts: MIT Press.
- Crane, T. 1995. *The Mechanical Mind: A Philosophical introduction to minds, machines and mental representation*. London: Penguin.
- Damasio, Antonio. 1996. *Descartes' Error: Emotion, Reason and the Human Brain*. London: Papermac-Macmillan.
- Dennet, D. 1998. *Brainchildren: Essays on Designing Minds*. Cambridge: MIT Press.
- Dyson, G. 1997. *Darwin among the Machines*. Massachusetts: Addison-Wesley.
- Franklin, Stan. 1995. *Artificial Minds*. Cambridge, Massachusetts: MIT Press.
- Gelernter, D. 1994. *The Muse in the Machine: Computerizing the Poetry of Human Thought*. New York: Free Press, Macmillan.
- Gershenfeld, N.A. 1999. *When Things Start to Think*. New York: Henry Holt & Co.
- Harle, R.F. 1999. *Philosophy and Psychoanalysis: The Myth of the Unconscious*. Unpublished.
- Hobson, J.A. 1988. *The Dreaming Brain*. New York: Basic Books.
- . 1994. *The Chemistry of Conscious States*. Boston: Little Brown & Co.
- Hofstadter, D.R. & Dennett, D.C. 1981. *The Mind's I: Fantasies and Reflections on Self and Soul*. Middlesex: Penguin.

- Hofstadter, D.R. 1995. *Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought*. New York: Basic Books.
- Kurzweil, Ray. 1999. *The Age of Spiritual Machines: When Computers Exceed Human Intelligence*. New York: Viking.
- Minsky, M. 1987. *The Society of Mind*. London: Heinemann.
- Moravec, Hans P. 1999. *Robot: Mere Machine or Transcendent Mind*. New York: Oxford University Press.
- Partridge, Derek. 1998. *Artificial Intelligence and Software Engineering: Understanding the Promise of the Future*. New York: Amacom.
- Searle, J.R. 1983. *Intentionality: An Essay in the Philosophy of Mind*. Cambridge: Cambridge University Press.
- . 1992. *The Rediscovery of the Mind*. Cambridge, Massachusetts: MIT Press.
- . 1995. *The Construction of Social Reality*. New York: Free Press.
- Thelen, E. & Smith, L.B. 1994. *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge, Massachusetts: MIT Press.
- Vinge, V. 1993. *Vision-21*. Symposium, NASA Lewis Research Centre & Ohio Aerospace Institute.
- Weizenbaum, J. 1976. *Computer Power and Human Reason*. London: Penguin.