

Review of Consciousness and the Possibility of Conscious Robots

Lyle N. Long*

Pennsylvania State University, University Park, Pennsylvania 16802

and

Troy D. Kelley†

U.S. Army Research Laboratory, Aberdeen, Maryland 20783

DOI: 10.2514/1.46188

This paper discusses the psychological, philosophical, and neurological definitions of consciousness and the prospects for the development of conscious machines or robots in the foreseeable future. Various definitions of consciousness are introduced and discussed within the different fields mentioned. A conscious machine or robot may be within the realm of engineering possibilities if current technological developments, especially Moore's law, continue at their current pace. Given the complexity of cognition and consciousness a hybrid parallel architecture with significant input/output appears to offer the best solution for the implementation of a complex system of systems which functionally approximates a human mind. Ideally, this architecture would include traditional symbolic representations as well as distributed representations which approximate the nonlinear dynamics seen in the human brain.

I. Introduction

WHILE there have been numerous discussions of computers reaching human levels of intelligence [1–3], building intelligent or conscious machines is still an enormously complicated task. Kurzweil [3] believes there will be systems with intelligence equal to humans by the late 2020s, and that we will see a merging of human and machine systems. Philosophers [4–6] and psychologists [7,8] have been debating consciousness for centuries, and more recently neuroscientists have begun discussing the scientific aspects of consciousness [9–13]. Discover Magazine [Nov. 1992] referred to consciousness as one of the “ten great unanswered questions of science.” It is time for engineers and scientists to seriously discuss the architectural requirements and possibilities of building conscious systems. This paper compares and contrasts what is known about consciousness from philosophy, psychology, and neuroscience with what might be possible to build using complex systems of computers, sensors, algorithms, and software. This paper has three purposes: 1) to review the current understanding of consciousness in a form suitable for engineers, 2) to discuss the possibility of conscious robots, and 3) to give some preliminary architectural requirements for conscious robot designs.

II. Definitions: Autonomy, Intelligence, and Consciousness

It is important to distinguish between autonomy, intelligence, and consciousness. In the field of unmanned vehicles (air-, land-, or sea-based) the terms autonomous and intelligent are often used synonymously, but these are different

Received 29 June 2009; accepted for publication 29 November 2009. Copyright © 2009 by Lyle N. Long and Troy D. Kelley. Published by the American Institute of Aeronautics and Astronautics, Inc., with permission. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 1542-9423/09 \$10.00 in correspondence with the CCC.

* Distinguished Professor, Aerospace Engineering, Bioengineering, and Mathematics, AIAA Fellow, lnl@psu.edu.

† Engineering Psychologist, Human Research and Engineering Directorate.

ideas. Many unmanned systems are simply operated remotely, however, they do not have any onboard intelligence. Intelligence can be defined as [14]:

“A very general mental capability that, among other things, involves the ability to reason, plan, solve problems, think abstractly, comprehend complex ideas, learn quickly and learn from experience.”

Autonomy is different than intelligence and consciousness [15]:

“Autonomy refers to systems capable of operating in the real-world environment without any form of external control for extended periods of time.”

A system can be autonomous, but not very intelligent (e.g., an earthworm) or it could be intelligent but not autonomous (e.g., a supercomputer, with appropriate software to simulate intelligence). Autonomy would require some intelligence, however. Clearly, it is possible to have varying levels of both autonomy and intelligence. It is also possible to have varying levels of consciousness.

Intelligence and consciousness are not the same thing either, and will have different architectural requirements. A conscious system would have capabilities far beyond a mere intelligent or autonomous system, but one of the problems in the scientific study of consciousness is that people often interpret consciousness differently. In some cases, people will take it to mean something far beyond self-awareness. While not everyone agrees on a definition of consciousness, one well-accepted definition [12] describes it as a “state of awareness” or being self-aware including:

- 1) *Subjectivity*: Our own ideas, moods, and sensations are experienced directly; unlike those of other people.
- 2) *Unity*: All sensor modalities melded into one experience.
- 3) *Intentionality*: Experiences have meaning beyond the current moment.

These arise simply from the physical properties of the neurons and synapses in the central nervous system [12], not some mystical properties (as Descartes claimed [5]) or quantum effects (as Penrose and others claim [16]). In addition, consciousness is often closely associated with attention [10,17]. Attention brings objects into our consciousness, and also allows us to handle the massive amounts of data entering our brains, however, some things are attended to unconsciously. Another definition of consciousness that is often cited is [18]:

“Most psychologists define consciousness simply as the experiencing of one’s own mental events in such a manner that one can report on them to others.”

The above two definitions are often called self-consciousness or access consciousness. Esoteric questions such as do humans all perceive the color red in the same manner, or what does it feel like to be a cat, or what it is like to be a particular cat [19] will not be considered here. Some say the big problem with consciousness is that there is no definitive test for it [3], so it is difficult to address scientifically, compounded by the problem that there are many different definitions of consciousness. If one restricts ourselves to testing if something or someone is self-aware or self-conscious, then there probably are tests. It is likely that machines can (and will) be self-aware, we can test for it, and it will be a remarkable moment in history.

While autonomy and intelligence are uncoupled, consciousness is related to intelligence and there are probably gradations in consciousness. Many people believe that many mammals have some level of consciousness or are at least self-aware; and there are even indications that fish may have consciousness [20–22]. One simple test for this is the “mirror test,” where a spot of color is placed on the test subject and when the subject looks in the mirror they recognize that they are seeing themselves (maybe by trying to touch the spot on their own body not the mirror). Humans older than 18 months, great apes, bottlenose dolphins, pigeons, elephants, and magpies all pass this test and show apparent self-awareness. When consciousness is defined as above, it is not that difficult to speculate that in the future machines will be conscious, but in order for them to have subjectivity, unity, and intentionality they will need powerful processing power, significant multi-model sensory input with data fusion, machine learning, and large memory systems.

One model of the varying levels of intelligence [23] (and probably consciousness) is shown in Fig. 1. At the lowest level are creatures that just perform stimulus-response behavior such as worms. Simple versions of robots such as these are fairly easy to build, for example, with simply a touch sensor and simple motor control. At the next level (e.g., a goldfish) there is significant perception, sensor input, and sensor processing. The structure of the

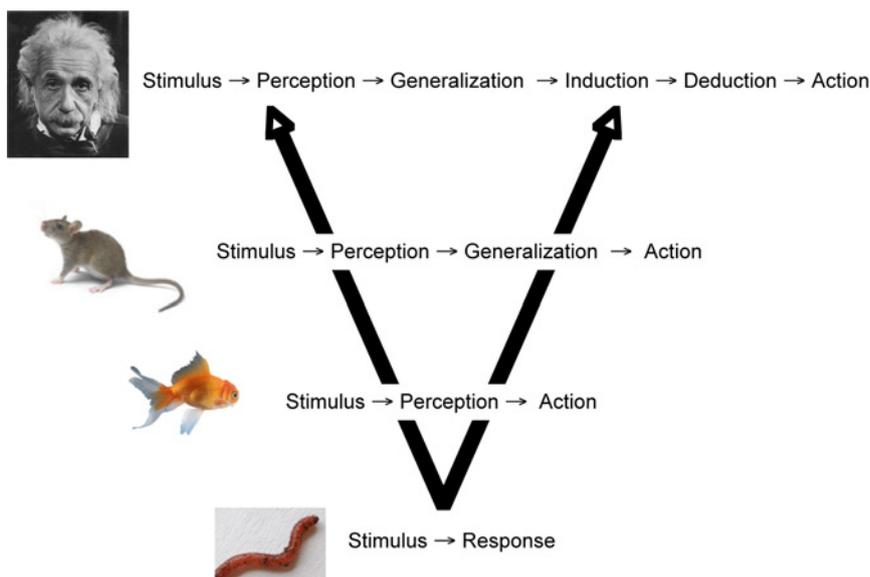


Fig. 1 Levels of intelligence.

goldfish brain has many similarities to mammalian brains [20]. At the next level (e.g., rats), the system is capable of generalizing its experience, i.e., applying its current knowledge to analogous new situations. At the highest level, humans are also capable of induction and deduction. Along with levels of intelligence, there are levels of consciousness. We know that some non-human animals do exhibit self-awareness (possibly even goldfish). These levels of consciousness are related to the increasing levels of intelligent functionality (perception, generalization, induction, and deduction). Current robotic vehicles or systems have probably not achieved the intelligence of a rat or the autonomy of a worm.

III. Views of Consciousness

Consciousness has been studied for about 2500 years, and in several different fields of study, including Philosophy, Psychology, Cognitive Science, Neuroscience, Computer Science, and Artificial Intelligence (AI). A brief review of some of the history, key figures, and key publications of consciousness will be presented, including the current understanding of the mammalian brain structure from neuroscience.

A. Philosophy

Philosophers have been debating consciousness and its implications for thousands of years. Some of the earliest studies related to consciousness are from Aristotle (384–322 B.C.) [24] in his discussions of the soul. Descartes (1596–1650) is known for his famous quote “Cogito, ergo sum” (“I think, therefore I am”). Descartes’ dualism maintains that the mind and the body are two different things, where the mind is immaterial and the body is material, and they somehow interact. It is remarkable that this concept is still being discussed today, when so much is known about neuroscience. Of course there are people who believe all sorts of things that make no logical sense or they do not believe things that have been scientifically shown to be true, as discussed by Dennett [25], Dawkins [26], and Hood [27]. Hood [27] says that “... adult supernaturalism is the residue of childhood misconceptions that have never been truly disposed of.” Consciousness will be explained completely once we understand more about the brain and nervous system (i.e., materialism and naturalism). Locke (1632–1704) [28], who was probably the first person to discuss consciousness as we define it today, considered consciousness to be the ability of a person to “consider itself as itself” and said “Consciousness is the perception of what passes in a man’s own mind.” This is still a good definition.

More recently, there have been several books on consciousness by modern philosophers, including Dennett [4,6], Searle [29], and Block et al. [19]. Dennett’s idea of consciousness [4] can be represented by the quote:

“Human consciousness is itself a huge complex of memes that can best be understood as the operation of a “von Neumannesque” virtual machine implemented in the parallel architecture of a brain that was not designed for any such activities,” where a meme is a cultural idea or a symbol. This is somewhat similar to how a conscious machine should be constructed. Dennett [6] also discussed the possibility of conscious machines. Searle said about whether a machine can be conscious:

“We have known the answer to this question for a century. The brain is a machine. It is a conscious machine. The brain is a biological machine just as much as the heart and the liver. So of course some machines can think and be conscious. Your brain and mine, for example.”

In the well-known Chinese Room concept due to Searle [29], one person is in a room filled with all the facts about the Chinese language and another person (outside the room) asks for translations. This is a bad example of the limitations of symbolic systems. This example is extremely over-simplified and it does not acknowledge that the human brain is a complex nonlinear system of systems [3]. Searle’s argument is that a simple translation machine will not understand the true meaning of the words being translated, however, future conscious machines will not be simple translation devices. Moreover, they will have both perceptual mechanisms and symbolic mechanisms and those two systems will be tightly coupled in order to convey meaning to the machine. Searle’s argument is like restricting our discussions of conscious robots to systems built using only a computer with simple input/output functionality. A conscious system will be a complex system of systems, with perceptions bound to symbolic knowledge, not just a simple input/output mechanism.

It may be that humans are not intelligent enough to understand their own brains, but this should not be surprising or disheartening. There are many very complex engineered systems being built today that cannot be understood by the human brain, but we can still design them and build them using the tools of engineering and systems engineering. With the help of psychology and neuroscience, we should be able to reverse engineer the human brain and central nervous system.

In the philosophy of mind there is also the notion of the “easy” and “hard” problems [19], and “access consciousness” and “phenomenal consciousness.” The “easy” cognitive problem (which is not necessarily easy at all!) refers to access consciousness and “discriminatory abilities, reportability of mental states, the focus of attention, and the control of behavior.” Most philosophers agree that these will eventually be explained via neuroscience. On the other hand, the “hard” problem refers to phenomenal consciousness or subjective experiences, e.g., “what it is like to be”. For example, Nagel [30] discusses how we cannot know what it is really like to be a bat, which has very different sensor input and processing than humans. For those interested in building conscious machines, the “hard problem” or “how it feels to be a ____” is not important in order to build conscious robots. If robots can be built that simply have self-awareness (which is agreed by most to be possible), it would be quite exciting.

Seager [31] says the mainstream view in philosophy is now that due to evolution:

“Consciousness *emerged* as a product of increasing biological complexity, from non-conscious precursors composed of non-conscious components.”

Philosophers agree on very little, but the majority agree that conscious (self-aware) machines are possible and will eventually be built.

This paper will deal with the more practical and useful “easy problem” and “access consciousness.” The “easy problem” can be explained by science and those systems can be reverse engineered, and therefore conscious (self-aware) intelligent machines can be built. Questions about what it feels like to be a cat or to sense the color red will be left to the philosophers, most of whom will be of little help in building conscious systems.

B. Psychology

William James (1842–1910) and Sigmund Freud (1856–1939) were two of the most influential psychologists. While Freud is more widely known by the general public, James was probably more influential. James [8] discussed consciousness a great deal. Two things he said in his book are:

“For practical purposes, nevertheless, and limiting the meaning of the word consciousness to the personal self of the individual, we can pretty confidently answer the question prefixed to this paragraph by saying that the cortex is the sole organ of consciousness in man.”

and

“My final conclusion, then, about the substantial Soul is that it explains nothing and guarantees nothing. I therefore feel entirely free to discard the word Soul from the rest of this book.”

Hood [27] says “When our body dies, so does our mind.” and “In short, misconceiving the mind lays the foundation for many of the beliefs in both religious and secular supernaturalism.”

Freud et al. [32] said: “The process of something becoming conscious is above all linked with the perceptions, which our sense organs receive from the external world.” Freud and his colleague, Pierre Janet, extensively studied hysterical amnesia brought about from traumatic events. They found that women who had been through trauma (death of a loved one) had sometimes developed the inability to remember events surrounding the trauma, and sometimes would completely forget the traumatic event itself. One woman had no recollection of an event where a man had mistakenly told her that her husband had died; yet the trauma was relived whenever the woman past the door where she was told the unfortunate news. The memory of the event was retained, but it was no longer available for her conscious inspection. Freud went on to develop seminal theories in psychotherapy based on his work on hysterical amnesia. His main contribution to the field being that unconscious memories of past traumatic events could have a powerful effect on conscious behavior. Gray [18] includes an excellent modern discussion of our current understanding of consciousness from a psychological perspective.

The American psychologist, Morton Prince, was one of the first to delineate unconscious behavior in normal people from the pathological women whom Freud and Janet had previously studied [33]. Cognitive psychologists eventually came to discover incidences where improvements or deficits in learning and memory were not explicitly recalled by patients. These memories and learning events were later to be defined as “implicit memories”. These were memories that were apparent in behavior, however, the memories were not immediately accessible to conscious experience. As Schacter put it, “a memory for a recent event can be expressed explicitly, as a conscious recollection, or implicitly, as a facilitation of test performance without conscious recollection”.

Neurological psychologists also discovered unusual memory deficits, which were first identified by Sergei Korsacoff, as amnesia. Korsacoff’s syndrome, as it was later to become known, described a syndrome in which subjects frequently had no recollection of events leading up to, or immediately after, certain events. In other words, amnesiacs displayed complete recall for certain memories, but complete lack of recall for other memories. Psychologists would later conclude that amnesia was a physiological indication of the separation between memory components and the apparent multiple functionality of the memory system, however, the exact nature of this separation continues to be debated. Nevertheless, the dual mechanism theory of memory would later emerge. Memories would be categorized as either procedural or declarative. Procedural memory was unconscious, episodic, and instance-based while declarative memory was more conscious, semantic, and fact based. This type of separation also maps well onto the symbolic and subsymbolic distinctions of knowledge organization [34]. Procedural memory is easily characterized as subsymbolic, while declarative memory is more easily characterized as symbolic.

As the theory has developed there have also been debates as to whether or not implicit memory can be represented at all within a symbolic system. For example, Cleeremans [35], argued that implicit learning did not fit well within the traditional cognitive science framework put forth by Newell and Simon [36], and that implicit learning was not a reasonable learning mechanism within a symbolic framework. However, symbolic proponents argued that implicit learning could occur symbolically, and in fact it was the same as the symbolic/declarative system, but simply lacked conscious awareness. Additionally, symbolic proponents argued that procedural memories were simply instance-based memories which could easily be represented as rules. These arguments and debates continue within the symbolic and subsymbolic knowledge representation communities.

Irrespective of representation, there appears to be some cognitive advantages to unconscious or implicit memories. For example, Jacoby and Dallas [37] found that priming effects (procedural memories) were persistent as memories for days and weeks while recognition memory (declarative memories) were not. This difference in retention rate could have certain evolutionary advantages. This would be beneficial to many cognitive systems, especially if the memories persisted for a long time, and were relatively easy to assimilate, given an adequate number of synapses. Simple unconscious associations of environments (the watering hole) to possible threats (lions) make it an attractive learning mechanism for many animals. Later research indicated that there was a certain amount of task dependency in priming effects, and that under certain conditions, the priming effects were greatly diminished [38]. However, the

discovery of task dependencies is typical in psychology and does not diminish the discovery that priming effects under certain circumstances persist longer than declarative memories. These long lasting, unconscious memories, certainly would be advantageous to early human beings.

Psychology has also taught us of the importance of the unconscious, in some ways the unconscious processes are even more interesting and mysterious than the conscious. Baars [39] discusses the importance of conscious and unconscious processes in describing his global workspace system, which has recently seen some possible experimental support [40]. Lehrer [41] has many great examples of the importance of the unconscious in human behavior. In some instances humans need to make decisions based on emotion, while in other cases they should rely on rational thought. Effective conscious systems will need this same capability. People often get into trouble by using the wrong approach. And usually it is only relatively simple problems that we are capable of solving rationally due to the limitations of the prefrontal cortex. The unconscious portions of the brain can be trained, and to become an expert in most things (music, sports, driving, etc.) usually requires 10,000 h of practice [42]. It is the “magic number of greatness.” This is roughly five years of full-time work, which also corresponds to the time it takes to get professional degrees (engineering, architecture, etc.) and the time it takes to get a Ph.D. During this period the brain is being trained to act quickly without long deliberations. For example, a major league baseball batter does not have time to think about his swing after the pitch, he just has to do it—and do it fast. He does it using what Koch [10] calls “zombie agents” in his brain (sometimes inaccurately called “muscle memory”).

The notions of conscious and unconscious processes are important to consider for robots as well. Essentially, all the processing that occurs in existing robots is unconscious. If a robot becomes conscious, will it also exhibit emotions? Robots will most likely be more capable if they can balance emotional and rational thinking, and have a mix of conscious and unconscious processing.

Emotions, such as love, hate, fear, sadness, anger, and remorse, have evolved in humans to help them survive. Abilities such as these will help conscious machines as well. In fact, it may be that emotions only occur in conscious systems. There have been several recent studies [43] of using cognitive architectures to model emotions. In fact, if machines can be imbued with emotions, they may also experience feelings such as depression.

One also needs to differentiate between the process of *developing* consciousness (fetal and human development) and the systems required to *maintain* consciousness. In trying to engineer conscious systems, the latter is much less important. We really need to understand how systems grow, learn, and become conscious in order to build conscious machines. This is discussed in more detail later in this paper.

C. Neuroscience

The most well-known and complete reference to neuroscience is the book by Kandel et al. [12], but the books by Churchland and Sejnowski [11], Koch [10,44], Pinker [7], Kandel [45], LeDoux [46], and Lehrer [41] are fascinating works. Crick and Koch [47] discuss the possible neural correlates of consciousness, but this issue is still unresolved.

The human brain, which is the most complicated system in the known universe, has been evolving for at least 4 million years. It uses roughly 20% of the energy in the human body. Evolution is basically an optimization program, and there are good reasons why the brain has evolved to its present state, as described in the books by Pinker [7], LeDoux [46], and Dennett [25]. LeVay [48] said “The mind is just the brain doing its job.” Genetic algorithms and evolutionary techniques could be used to simulate human evolution; however, duplicating the conditions that led to the evolution of the human brain would be difficult, if not impossible [7,25].

The brain, in engineering terms, is a complex nonlinear system of systems. These include:

- Cerebellum
 - The “little brain” or reptilian brain
 - Motor movement
 - If this is not functioning, you will not be conscious
- Cerebral cortex
 - Newer part of the brain for mammals
 - Plays a major role in consciousness
- Perceptual cortices
 - Visual cortex
 - You do not see with your eyes, you see with your brain

- Visual processing can take up to 25% of the brain
- Auditory cortex
- Somatosensory system
- Limbic system
 - Memory system
 - Olfactory system
 - Emotions and memories get tied together here

These systems do perform parallel computing, i.e., the brain does many things simultaneously. And they are all highly inter-connected.

In addition to the processing systems, there are innumerable interconnection networks. The white matter in the brain represents pathways between the different areas of the brain (the gray matter represents neurons). The neural pathways take three forms:

- Association pathways (within cortex)
- Commissural pathways
 - Corpus collosum has 300 million axonal connections
- Projection pathways
 - Connections between brain subsystems (e.g., motor cortex to muscles)

The cortex is the area of the brain that does most of the higher-level processing, and is the newest (evolutionary speaking) portion of the brain. It is the seat of the mind [49]. The cortex is a wrinkled sheet on the outer edge of the brain. It is about 0.2 m² and about 1–4 mm thick. The neocortex has six layers, older parts of the cortex have only three layers.

The cortex is just one part of the brain, but it has about 50–200 different regions within it, each a separate subsystem. Vision alone uses roughly 40 different regions of the cortex. Brodmann divided the cortex into 50 regions [12]; for example Area 17 is the primary visual cortex. The cortex has four main lobes (each of which has subregions):

- Frontal lobe
 - Planning, higher-cognitive functions
- Parietal cortex
 - Somasensory, written language, vision and somasensory data fusion
- Occipital lobe
 - Primary visual cortex
- Temporal lobe
 - Auditory cortex, face recognition, hippocampus, learning and memory

Clearly, the brain is a hierarchical system of systems, and humans have begun to understand many parts of it. In the not too distant future, it should be possible to reverse engineer much of it.

The subsystems of the brain and CNS are not logic (symbolic) processing units. Nor are they floating point or integer processors as in a computer. They are spiking neural networks [50–52] with neurons that fire at roughly 50 Hz. In addition, there are roughly 150 different kinds of neurons in the human body. Table 1 shows the brain weight and number of neurons for several different animals [53]. Humans have roughly 12 billion cortical neurons (5 billion more than chimpanzees), while the entire brain has about 100 billion. The human cerebral cortex has roughly 10¹⁴ synapses. If you consider the memory potential of a single synapse as 1 byte, then the human cortex has roughly

Table 1 Brain weight and number of cortical neurons

| Animal | Brain weight (gm) | Number of neurons |
|------------|-------------------|-------------------|
| Human | 1350 | 100,000,000,000 |
| Chimpanzee | 420 | 30,000,000,000 |
| Cat | 30 | 300,000,000 |
| Rat | 2 | 15,000,000 |
| Mouse | 0.3 | 4,000,000 |
| Zebrafish | ~ 0.001 | 10,000 |
| C. Elegans | — | 300 |

10^{14} bytes (100 terabytes). Another estimate, by Kurzweil [3], is that humans can store 10^7 “chunks” of knowledge, and he says a chunk is roughly 10^5 bytes (or 10^{12} bytes total). A conscious machine equal to humans will need this level of processing power, learning ability, memory capacity, and network interconnections. So a machine that uses a symbolic approach might need to store roughly 10^7 chunks, while a connectionist approach might require 10^{12} – 10^{14} synapses and 10^{11} neurons.

Computational models of neurons have been developed. In order of decreasing complexity, some of the models are: Hodgkin–Huxley (H–H) [54], Fitzhugh [55], Izhikevich [56], and leaky integrate and fire (LIF) [44]. The H–H model involves four nonlinear coupled ordinary differential equations (ODE) for each neuron, and is very expensive to compute. The LIF model is a single nonlinear ODE for each neuron, and only requires roughly 1 byte of storage per neuron and one floating-point operation per neuron per time step [50]. Gupta and Long [57] have used LIF combined with Hebbian learning to simulate portions of the mammalian vision system.

Penfield [58] performed some very interesting studies while performing neurosurgery. He directly stimulated portions of the human brain and the patients experienced movements, tastes, smells, etc. They also vividly relived past memories. This clearly demonstrated that the mind is a product of the physical brain. Even the emotion of love can be traced to the neurotransmitter oxytocin. While the brain is the seat of the mind, the brain is just one part of the central nervous system. An adult human might still be conscious even if they lose these systems, but it is unlikely that a newborn could become conscious without them.

In discussing conscious systems, it is important to differentiate between what it takes to *become* conscious and what it takes to *maintain* consciousness. Koch and Tononi [13] discuss the possibility of conscious machines, but most of their descriptions concern what a conscious system needs, or does not need, to remain conscious; and this applies to systems that are already conscious. That is, if portions of an adult brain were removed, would the adult human still be conscious—well, it depends on which portions are removed. The well-known physicist Stephen Hawking (who has Lou Gehrig’s disease) is obviously conscious, but if his condition had developed before he became conscious, he might never have achieved consciousness. This could also be said of Bauby [59], described in the incredible book and film called *The Diving Bell and Butterfly*. In developing conscious machines, it is most important to consider how does something *become* conscious. This is because it will not be possible to simply build a conscious machine and program every aspect of it. Building a conscious system is more like a developing embryo or teaching an infant, both of which eventually become conscious. One needs to consider what functions are critical in the *development* of consciousness and how one can build a system that will learn and experience enough to become conscious.

In addition, because the human nervous system is so complicated, it might be extremely worthwhile to study simpler systems in animals. Mammals have been studied extensively, but even simpler systems would be valuable to study. For example, fish have brains and nervous systems that have surprising similarities to mammals, and they are much simpler to study. Once these simpler systems are understood, they can be scaled to larger systems.

Given the appropriate level of computing hardware, algorithms, sensors, learning ability, and memory, the machine will have to experiment and explore its environment. As it learns about its environment, it will also learn about itself, and consciousness may emerge. A fetus/newborn becomes conscious somewhere between the third trimester and fourth year, depending on how you define consciousness. Learning is essential because it will not be possible to program the entire machine. It needs to learn like a human infant would learn and then it may recognize that it exists, in addition to all the other objects it learns about. It would also be very useful if there were humans there to teach it, as they teach infants. It is important to remember, however, that it takes a very long time for humans to learn to function as an adult (roughly 18 years).

D. Computer Science and AI

There have been many comparisons between computers and brains [1,2], but these are very different systems. Philosophers often make the mistake of comparing the human brain to a “computer.” There are some very small portions of the brain that might be compared to a computer. For example, maybe the retina could be compared to a computer, because it has fairly well-known input/output channels and does some fairly simple processing, but the brain is not like a typical computer at all. First of all, the brain is an analog device, not a digital device. Even the vision system is a very complex system [60], and more complicated than a typical computer. In reality the human

brain is thousands of parallel and inter-connected neural networks with many parallel channels of input (sensory neurons) and output (motor neurons). As mentioned earlier, Dennett says [4]:

“Human consciousness is itself a huge complex of memes that can best be understood as the operation of a “von Neumannesque” virtual machine implemented in the parallel architecture of a brain that was not designed for any such activities.”

In engineering terms:

Human consciousness comes about from a highly interconnected complex system of systems using nonlinear spiking neural networks to perform data fusion on vast amounts of input data to learn, to store memories, to think, and to control a complex motor subsystem.

Simple philosophical thought experiments about a single computer and single input/output channels are unlikely to provide useful conclusions about the brain and its subsystems. Logic alone cannot unravel the mysteries of the brain because it is too complex. The human brain and central nervous system are what engineers would call an extremely complex nonlinear system of systems.

Even the most powerful current supercomputer, the IBM RoadRunner [61] with 6,562 dual-core AMD Opteron chips and 12,240 Cell chips, cannot compare to the human brain. The RoadRunner has 9.8×10^{13} bytes of memory (98 terabytes) and could sustain 10^{15} operations per second (1 petaflop) peak speed. However, it is difficult to actually use the entire system at once and have code that is perfectly scalable. Also, the RoadRunner computer cost about \$100 M, requires 1000 m³ of space, weighs 228,000 kg, and requires 3.9 MW (while the human brain requires about 0.0014 m³ of space, weighs 1.3 kg, and needs 20 W). Figures 2 and 3 summarize this information on the human brain and this large supercomputer. While the supercomputer is close to being similar to the human brain in terms of memory and speed, the cost, size, weight, and power required are about 5–7 orders of magnitude worse than the brain.

Biological systems and computers can be compared in terms of memory and speed, but these are only two of the requirements for an intelligent machine. A typical laptop today has roughly the computational power of a cockroach nervous system [52], which is about a million times less powerful than the human brain. Assuming Moore’s law remained valid, it would take about 40 years for laptops to reach human computational power.

While the brain uses neurons and synapses, these alone will not lead to conscious machines using current computers, so in the near term hybrid systems will be required. Also, in trying to duplicate the human brain in an artificial system, there are many other unknowns such as neuronal interconnections (wiring diagrams), hardware, algorithms, learning, sensory input, and motor-control output. There is an enormous need for the development of processing hardware more efficient than traditional off-the-shelf processors for spiking neural networks. The DARPA SyNAPSE program [62] is one such attempt at that.

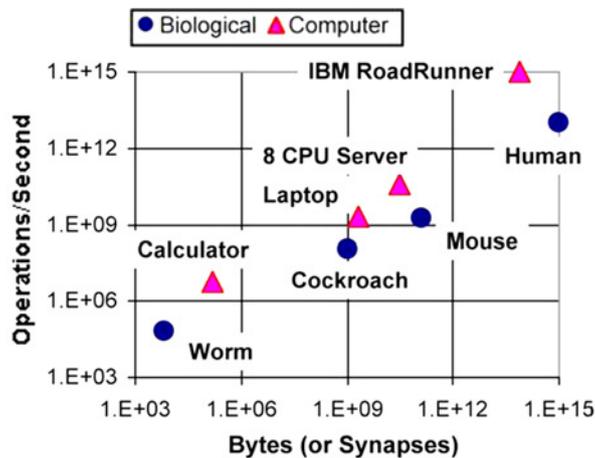


Fig. 2 Processing power and memory capabilities of biological and manmade systems.

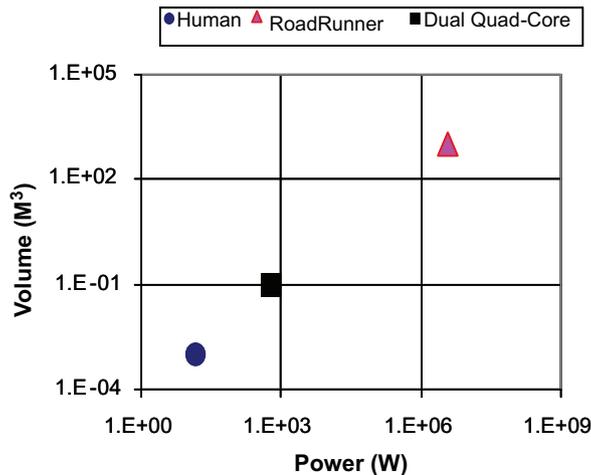


Fig. 3 Power and volume requirements for human brain, a server, and a supercomputer.

The field of AI has not produced very intelligent systems, because it has been too focused on symbolic processing. Even cognitive architectures (e.g., [63], executive-process/interactive control (EPIC) [64], state operator and result (Soar) [65,66], symbolic and subsymbolic robotic intelligence control system (SS-RICS) [66,67], and adaptive control of thought—rational (ACT-R) [68,69]), which have been implemented on mobile robots [66,70], are not even close to human intelligence and power and have only rudimentary learning ability. Engineers involved in computational intelligence are more focused on subsymbolic processes such as neural networks, genetic algorithms, and fuzzy logic, which may lead to large-scale intelligent systems. The most effective approach for the near term, however, will be hybrid methods that combine symbolic and subsymbolic approaches.

IV. An Engineering Approach to Developing Conscious Systems

The brain is a complex nonlinear system of systems, and engineers are quite capable of building such systems. Minsky [71] said:

“There is not the slightest reason to doubt that brains are anything other than machines with enormous numbers of parts that work in perfect accord with physical laws. As far as anyone can tell, our minds are merely complex processes. The serious problems come from our having had so little experience with machines of such complexity that we are not yet prepared to think effectively about them.”

But humans *do* now have experience with machines of enormous complexity. Consider:

- Microsoft Windows Vista operating system: 50 million lines of software
- Boeing 777 aircraft: 3 million lines of software, 1100 processors, and 3 million parts
- Intel Tukwila Chip: 2 billion transistors
- Internet: 4 billion addresses (using IPv4) and 10^{38} addresses (using IPv6)
- IBM RoadRunner Supercomputer: 6,562 dual-core AMD Opteron chips and 12,240 IBM Cell chips

The use of well-established engineering approaches (especially Systems Engineering and Software Engineering [72]) allow the successful completion of very complex engineering projects. No single human can fully comprehend a Boeing 777 or the Internet, but they can be built. In addition, the brain is not made up of that many different parts. There are roughly only 150 different kinds of neurons. Humans can certainly build systems with billions of parts if the parts are quite similar.

A conscious machine cannot just be an isolated computer. It will need to be a complex system of systems with an enormous amount of sensor data, and it must be capable of learning and understanding real-world situations. The key, however, will be emergent behavior development through a variety of algorithmic techniques including for example: genetic algorithms, machine learning, fuzzy logic, cognitive architectures, and neural networks. Humans will not be capable of completely specifying and programming the entire system; learning and emergent behavior [10] will be a stringent requirement for development of the system.

Conscious machines will also need to be embedded in the real world with significant input/output capabilities and the ability to learn from people and experience. Robots learning from humans, however, require care and dedication on the part of the humans. The sharing of knowledge is difficult and costly for the one doing the sharing—as every parent knows. Human infants are at an evolutionary advantage because their parents have an evolutionary stake in the development of their children. Robots do not have this evolutionary stake. So, human–robotic interaction needs to be crafted to keep the interest of the human high and to offer some benefits and advantages to the human who is offering the knowledge to the robot.

One particularly interesting group of robots are the humanoid and quadruped robots that have been developed recently [73–77]. The Honda and Sony Corporations had some very interesting robots [73,74,77], but in 2006 Sony decided to cease their activities in this area. They made great progress in gaits, dynamics, and control. An MIT group [75] constructed “COG,” which had 21 degrees of freedom and many sensors (visual, auditory, vestibular, kinesthetic, and tactile). They discuss the need to move away from classic AI methods such as monolithic models/control and general purpose computing. They also discuss the need for robots to develop similarly to how humans develop. The human–machine interactions of humanoid and quadruped robots will also help us understand both robots and humans.

The human sensory systems use hundreds of millions of cells, and there are roughly 600 muscles in the human body. The fascinating robotic vehicles in the DARPA Urban Challenge have very few sensor systems (e.g., lasers, cameras, and radar) and very few motor-control output channels. They also required millions of lines of software and teams of engineers, and they still have little or no learning ability.

In the near term, we will need hybrid systems: symbolic and subsymbolic, (e.g., cognitive architectures and neural networks will both be important). While the human brain uses neurons for processing, it also uses collections of neurons to perform rule-based processing similar to cognitive architectures. It is important to concentrate on the functional aspects of cognition (i.e., learning, memory, decision making) which are well represented in today’s cognitive architectures. Attempting to duplicate entire neurological systems (i.e., the thalamus) is useful for understanding neurological and anatomical relationships within the brain; however, this might be too labor intensive and it is probably unnecessary for creating a conscious system. Rather, concentrating on the functional aspects of cognition (i.e., the thalamus is a gateway processor for sensory input) will lead to greater replication of higher-level cognition than the reproduction of entire neurological systems. Additionally, consciousness will result as an emergent behavior if there is adequate sensor input, processing power, and robust learning. Current robotic systems are many orders of magnitude away from human abilities.

Emergence is a crucial property. Biological consciousness emerged historically due to evolution. Consciousness is also an emergent property of every human. While humans are born with enormous sensory processing abilities,

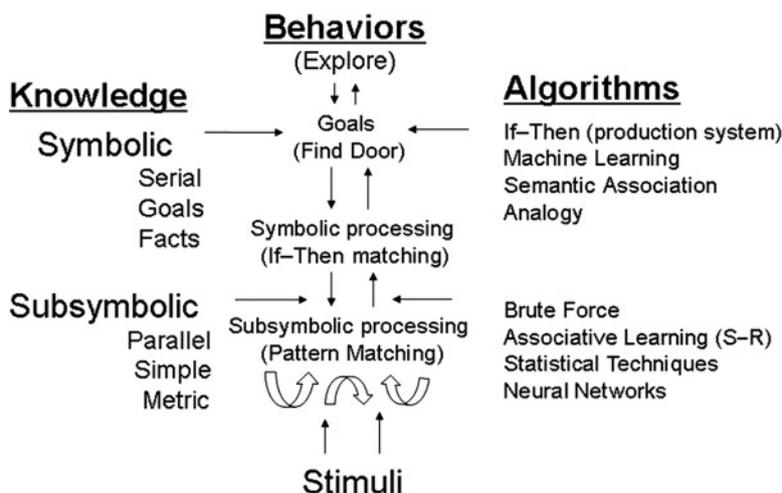


Fig. 4 SS-RICS notional data flow emphasizing the symbolic and sub-symbolic distinctions [51].

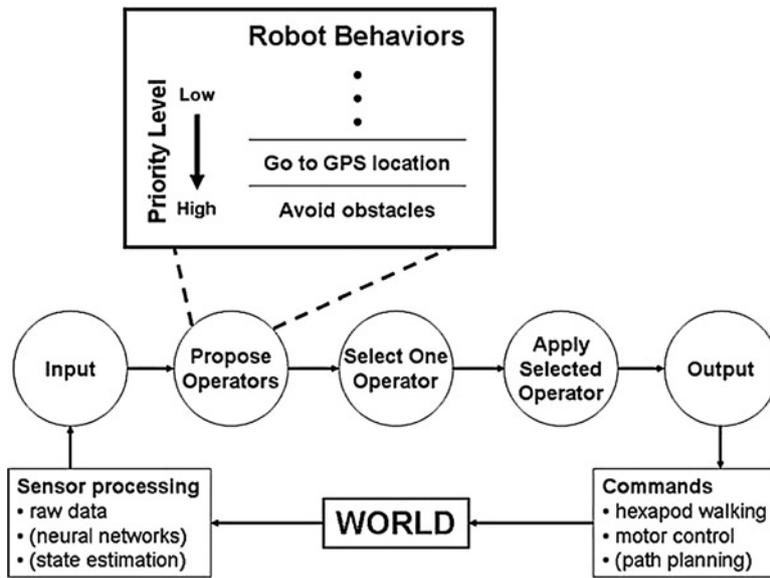


Fig. 5 Soar cognitive architecture implemented on a mobile robot [50].

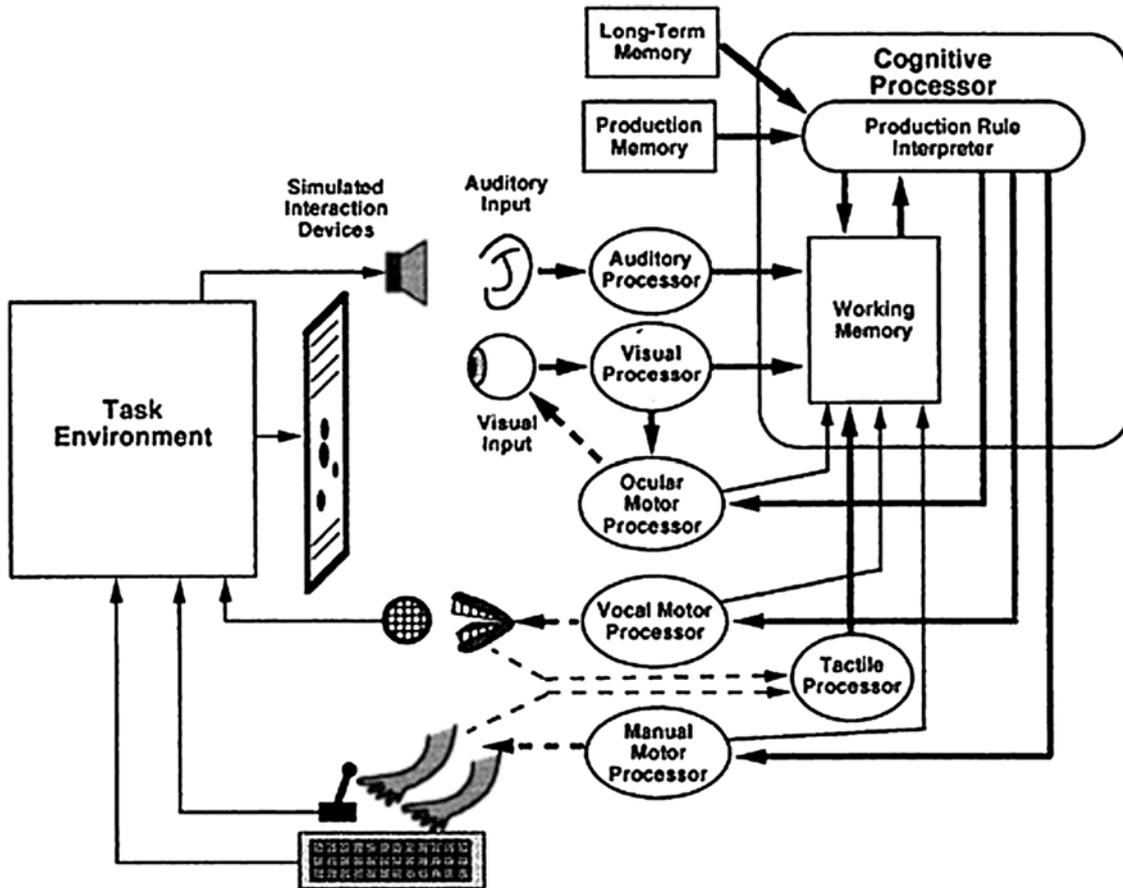


Fig. 6 EPIC cognitive flowchart [68].

enormous neural networks, and sophisticated motor control systems; they are unlikely to be conscious or self-aware at birth. Only after learning and interacting with its environment due humans become fully conscious. This is a very important point because very sophisticated robots can be built, but we will need to rely on the emergent property of consciousness for the robots to become self-aware.

The architecture of the SS-RICS [66,67] shown in Fig. 4 is one attempt at a general-purpose intelligent system. This system uses a cognitive architecture (based on the ACT-R [69]) for decision making, but allows for subsymbolic processing, such as neural networks, for processing sensor data. The Soar cognitive architecture has been coupled to sensor inputs and motor outputs [70], as shown in Fig. 5. The sensor data is processed subsymbolically and then symbols or states are passed to Soar. Figure 6 shows a flowchart for the EPIC cognitive architecture from Meyer and Kieras [64]. This too shows how sensor data can be fed into the cognitive architecture. These hybrid approaches will be needed in the near term to build intelligent systems.

Figure 7 shows a more general flowchart and the importance of input data processing, data fusion, and memory. Humans have hundreds of millions of input sensors (e.g., visual, taste, olfactory, haptic, and auditory), and these signals are thoroughly processed through a complex hierarchical neural system. Sensor input and processing is highly parallel and involves significant learning. Conscious robots will need similar levels of sensor input and sensor processing, but they will not be limited to the same sensors as humans, e.g., they could have infrared sensors, electromagnetic sensors, compasses, GPS, etc. In addition to the high levels of sensor input and processing, humans have (and robots will need) the ability to store subsymbolic information and to perform sensor data fusion. It is much easier and more accurate to recognize an object (for example a cat) by combining what we know about how it looks, sounds, feels, and smells. Three other important functions for humans and robots are motor control, symbolic memories, and attention, which will need to be tied to a cognitive system.

As shown in Fig. 1, low-level biological systems (e.g., goldfish) can perform sensing, perception, and action, but to achieve consciousness the system will need to be able to perform generalization, induction, and deduction. Traditional AI has not achieved the lofty goals originally proposed because most AI systems cannot generalize. The IBM computer (Deep Blue) that beat the chess champion Kasparov can *only* play chess. It cannot play checkers. It cannot generalize

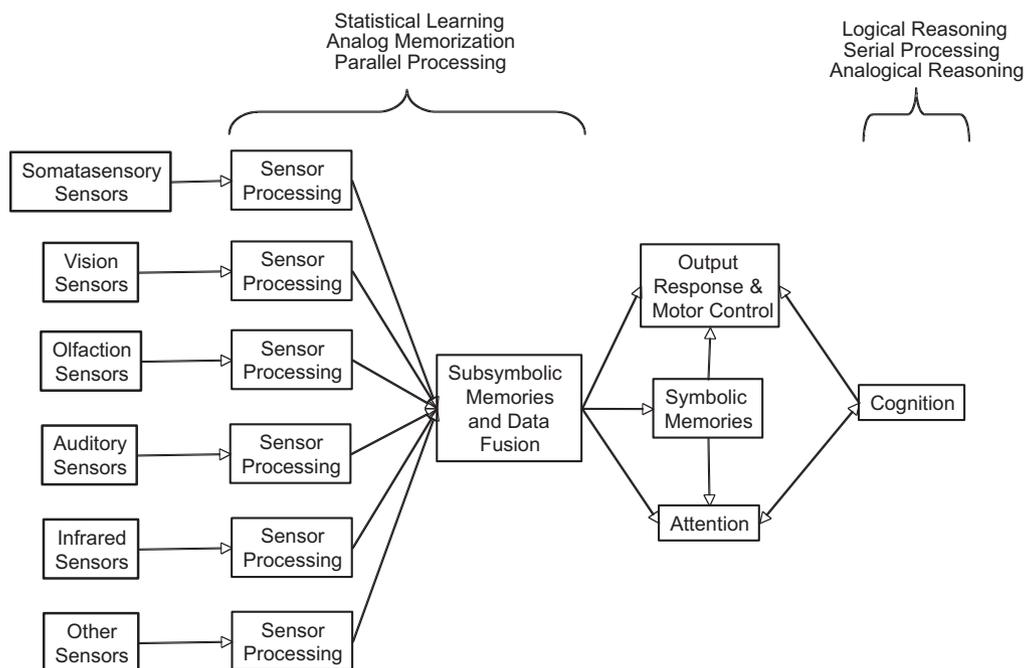


Fig. 7 Example architecture for future intelligent (possibly conscious) robots.

what it currently knows to new problems. Cognitive systems that can learn and perform generalization, induction, and deduction will be essential.

Symbolic AI alone will not lead to machines capable of duplicating human behavior. Connectionists and subsumptive architectures will not (in the near term), by themselves, lead to the development of human-level intelligence nor the functional characteristics that define consciousness. Rule-based systems and cognitive architectures that require humans to program the rules are not scalable to billions of rules (a.k.a the Frame problem [78]). The machines will need to rely on hybrid systems, learning, and emergent behavior; and they will need to be carefully taught and trained by teams of engineers and scientists.

V. Conclusion

Conscious robots are likely to be built within this century, if technological advances continue at their current pace. To accomplish this, an interdisciplinary effort involving neuroscience, psychology, computer science, and engineering will be needed. It will require a hybrid approach using functional approximations of cognition and this will probably include cognitive architectures, neural networks, fuzzy logic, data fusion, parallel computers, signal processing, and massive sensor arrays. It will also require computational power near the levels of the human brain. With proper cognitive development, hardware implementations, sensor processing, and efficient learning algorithms, machine consciousness is an achievable goal. It will not be possible, however, to simply build and program a conscious robot. Consciousness develops in humans as an emergent property, and we will need to rely on this for conscious robots as well. Conscious robots will have both positive and negative effects on society. In particular, ethics and laws defining what it means to be a human and what it means to be autonomous may need to be reconsidered; especially if there is a merging of human and machine systems. The first conscious robot that is built will be as astounding (and frightening) to humans as the discovery of life on other planets.

Acknowledgments

Lyle N. Long gratefully acknowledges support as a Moore Distinguished Scholar (2007–2008) at the California Institute of Technology, support from the Office of Naval Research (Grant No. N00014-05-1-0844), and support from the U.S. Army Research Laboratory (Contract No. TCN 07-305). We also acknowledge the valuable comments from Victoria Braithwaite and Scott Hanford. And finally, we would like to thank the reviewers for their insightful comments.

References

- [1] Kurzweil, R., *The Age of Spiritual Machines*, Penguin Putnam, New York, NY, 2000.
- [2] Moravec, H., *Robot: Mere Machine to Transcendent Mind*, Oxford Univ. Press, Oxford, 1998.
- [3] Kurzweil, R., *The Singularity is Near*, Penguin Group, London, 2005.
- [4] Dennett, D. C., *Consciousness Explained*, Back Bay Books, Boston, MA, 1992.
- [5] Descartes, R., *Meditations on First Philosophy*, Translated by John Veitch 1901, Reprinted by Prometheus Books, New York, NY, 1989.
- [6] Dennett, D. C., “The Practical Requirements for Making a Conscious Robot,” *Artificial Intelligence and the Mind*, Vol. A349, No. 1689, 1994, pp. 133–146.
- [7] Pinker, S., *How the Mind Works*, W. W. Norton, New York, NY, 1999.
- [8] James, W., *Principles of Psychology*, Holt, Rinehart, and Winston, New York, NY, 1890.
doi: [10.1037/10538-000](https://doi.org/10.1037/10538-000)
- [9] Edelman, G. M., *The Remembered Present: A Biological Theory of Consciousness*, Basic Books, New York, NY, 1989.
- [10] Koch, C., *The Quest for Consciousness*, Roberts & Co., Greenwood Village, CO, 2004.
- [11] Churchland, P. S., and Sejnowski, T. J., *The Computational Brain*, MIT Press, Cambridge, MA, 1992.
- [12] Kandel, E. R., Schwartz, J. H., and Jessell, T. M., *Principles of Neural Science*, McGraw-Hill Medical, New York, NY, 2008.
- [13] Koch, C., and Tononi, G., “Can Machines be Conscious?,” *IEEE Spectrum*, June 2008, pp. 55–59.
- [14] Gottfredson, L. S., “Mainstream Science on Intelligence: An Editorial with 52 Signatories, History, and Bibliography,” *Intelligence*, Vol. 24, No. 1, 1997, pp. 13–23.
doi: [10.1016/S0160-2896\(97\)90011-8](https://doi.org/10.1016/S0160-2896(97)90011-8)

- [15] Bekey, G. A., *Autonomous Robots: From Biological Inspiration to Implementation and Control*, MIT Press, Cambridge, MA, 2005.
- [16] Penrose, R., *Shadows of the Mind: A Search for the Missing Science of Consciousness*, Oxford Univ. Press, Oxford, 1996.
- [17] Koch, C., and Tsuchiya, N., "Attention and Consciousness: Two Distinct Brain Processes," *Trends in Cognitive Sciences*, Vol. 11, No. 1, 2007, pp. 16–22.
doi: [10.1016/j.tics.2006.10.012](https://doi.org/10.1016/j.tics.2006.10.012)
- [18] Gray, P., *Psychology*, Worth Pub., New York, NY, 2007.
- [19] Block, N., Flanagan, O., and Guzeldere, G., *The Nature of Consciousness: Philosophical Debates*, MIT Press, Cambridge, MA, 1997.
- [20] Brown, C., Laland, K., and Krause, J. (eds.), *Fish Cognition and Behavior*, Blackwell Science, Oxford, 2006.
- [21] Braithwaite, V. A., and Boulcott, P., "Can Fish Suffer?" in *Fish Welfare*, edited by Branson, E. J., Blackwell Publishing, Oxford, 2008.
- [22] Chandroo, K. P., Yue, S., and Moccia, D., "An Evaluation of Current Perspectives on Consciousness and Pain in Fishes," *Fish and Fisheries*, Vol. 5, No. 4, 2004, pp. 281–295.
doi: [10.1111/j.1467-2679.2004.00163.x](https://doi.org/10.1111/j.1467-2679.2004.00163.x)
- [23] Sowa, J. F. (ed.), *Categorization in Cognitive Computer Science*, Elsevier, Amsterdam, 2006.
- [24] Caston, V., "Aristotle on Consciousness," *Mind*, Vol. 111, No. 444, 2002, pp. 751–815.
doi: [10.1093/mind/111.444.751](https://doi.org/10.1093/mind/111.444.751)
- [25] Dennett, D. C., *Darwin's Dangerous Idea: Evolution and the Meanings of Life*, Simon and Schuster, Upper Saddle River, NJ, 1996.
- [26] Dawkins, R., *The God Delusion*, Houghton Mifflin, Boston, MA, 2006.
- [27] Hood, B. M., *Super Sense: Why We Believe in the Unbelievable*, Harper-Collins, New York, NY, 2009.
- [28] Locke, J., *An Essay Concerning Human Understanding*, T. Basset, London, 1690.
- [29] Searle, J. R., *The Mystery of Consciousness*, New York Review, New York, NY, 1990.
- [30] Nagel, T., "What is it Like to be a Bat?" *Philosophical Review*, Vol. 4, Oct. 1974, pp. 435–450.
doi: [10.2307/2183914](https://doi.org/10.2307/2183914)
- [31] Seager, W., "A Brief History of the Philosophical Problem of Consciousness," Chap. 2, edited by P. D. Zelazo, M. Moscovitch and E. Thompson, *The Cambridge Handbook of Consciousness*, Cambridge Univ. Press, Cambridge, MA, 2007.
- [32] Freud, S., Strachey, J., and Gay, P., *An Outline of Psycho-Analysis*, W. W. Norton, New York, NY, 1989.
- [33] Schacter, D. L., "On the Relation Between Memory and Consciousness: Dissociable Interactions and Conscious Experience," in *Varieties of Memory and Consciousness, Essays in Honour of Endel Tulving*, edited by H. L. Roediger and F. I. M. Craik, Erlbaum Pub., Hillsdale, NJ, Conference held in Toronto, May 1987.
- [34] Kelley, T. D., "Symbolic and Sub-symbolic Representations in Computational Models of Human Cognition: What Can be Learned from Biology?" *Theory & Psychology*, Vol. 13, No. 6, 2003, pp. 847–860.
doi: [10.1177/0959354303136005](https://doi.org/10.1177/0959354303136005)
- [35] Cleeremans, A., *How Implicit is Implicit Learning?*, Oxford Univ. Press, Oxford, 1996.
- [36] Newell, A., and Simon, H. A., *Human Problem Solving*, Prentice-Hall, Englewood Cliffs, NJ, 1972.
- [37] Jacoby, L. L., and Dallas, M., "On the Relationship Between Autobiographical Memory and Perceptual Learning," *Journal of Experimental Psychology*, Vol. 110, No. 3, 1981, pp. 306–340.
- [38] Cermak, L. S., Talbot, N., Chandler, K., and Wolbarst, L. R., "The Perceptual Priming Phenomenon in Amnesia," *Neuropsychologia*, Vol. 23, No. 5, 1985, pp. 615–622.
doi: [10.1016/0028-3932\(85\)90063-6](https://doi.org/10.1016/0028-3932(85)90063-6)
- [39] Baars, B. J., *A Cognitive Theory of Consciousness*, Cambridge Univ. Press, Cambridge, MA, 1993.
- [40] Gaillard, R., et al., "Converging Intracranial Markers of Conscious Access," *PLOS Biology*, Vol. 7, No. 3, 2009, pp. 472–492.
doi: [10.1371/journal.pbio.1000061](https://doi.org/10.1371/journal.pbio.1000061)
- [41] Lehrer, J., *How We Decide*, Houghton Mifflin, Boston, MA, 2009.
- [42] Gladwell, M., *Outliers: The Story of Success*, Little, Brown, and Co., Boston, MA, 2008.
- [43] Hudlicka, E., "Beyond Cognition: Modeling Emotion in Cognitive Architectures," *Proceedings of the International Conference on Cognitive Modeling (ICCM)*, Pittsburgh, PA, 2004, pp. 118–123.
- [44] Koch, C., *Biophysics of Computation: Information Processing in Single Neurons*, Oxford Univ. Press, Oxford, 1999.
- [45] Kandel, E. R., *In Search of Memory: The Emergence of a New Science of Mind*, W. W. Norton, New York, NY, 2006.
- [46] LeDoux, J., *Synaptic Self: How Our Brains Become Who We Are*, Penguin Putnam, New York, NY, 2003.
- [47] Crick, F., and Koch, C., "Towards a Neurobiological Theory of Consciousness," *Seminars in the Neurosciences*, Vol. 2, 1990, pp. 263–275.
- [48] LeVay, S., *The Sexual Brain*, MIT Press, Cambridge, MA, 1994.

- [49] Norden, J., *Understanding the Brain*, The Teaching Company, Chantilly, VA, 2007.
- [50] Long, L. N., "Scalable Biologically Inspired Neural Networks with Spike Time Based Learning," *Proceedings of the IEEE Symposium on Learning and Adaptive Behavior in Robotic Systems*, IEEE, Edinburgh, Scotland, 2008.
doi: [10.2514/1.31026](https://doi.org/10.2514/1.31026)
- [51] Long, L. N., and Gupta, A., "Scalable Massively Parallel Artificial Neural Networks," *Journal of Aerospace Computing, Information, and Communication*, Vol. 5, No. 1, 2008, pp. 3–15.
- [52] Long, L. N., and Gupta, A., "Biologically-Inspired Spiking Neural Networks with Hebbian Learning for Vision Processing," *46th AIAA Aerospace Sciences Meeting*, AIAA, Reno, NV, 2008, AIAA Paper 2008-0885.
- [53] Roth, G., and Dicke, U., "Evolution of the Brain and Intelligence," *Trends in Cognitive Sciences*, Vol. 9, No. 5, 2005, pp. 250–257.
doi: [10.1016/j.tics.2005.03.005](https://doi.org/10.1016/j.tics.2005.03.005)
- [54] Hodgkin, A. L., and Huxley, A. F., "A Quantitative Description of Ion Currents and Its Applications to Conduction and Excitation in Nerve Membranes," *Journal of Physiology*, Vol. 117, No. 4, 1952, pp. 500–544.
- [55] Fitzhugh, R., "Impulses and Physiological States in Models of Nerve Membrane," *Biophysical Journal*, Vol. 1, No. 6, 1961, pp. 445–466.
- [56] Izhikevich, E. M., "Which Model to Use for Cortical Spiking Neurons?" *IEEE Transactions on Neural Networks*, Vol. 15, No. 5, 2004, pp. 1063–1070.
doi: [10.1109/TNN.2004.832719](https://doi.org/10.1109/TNN.2004.832719)
- [57] Gupta, A., and Long, L. N., "Hebbian Learning with Winner Take All for Spiking Neural Networks," *Proceedings of the International Joint Conference on Neural Networks*, IEEE, Atlanta, GA, 2009.
- [58] Penfield, W., *The Mystery of the Mind: A Critical Study of Consciousness and the Human Brain*, Princeton Univ. Press, Princeton, NJ, 1975.
- [59] Bauby, J.-D., *The Diving Bell and Butterfly*, Alfred A. Knopf, New York, NY, 1997.
- [60] Martinez-Conde, S., and Macknik, S. L., "Windows on the Mind," *Scientific American*, Vol. 297, Aug. 2007, pp. 56–63.
doi: [10.1038/scientificamerican0807-56](https://doi.org/10.1038/scientificamerican0807-56)
- [61] TOP500 Computer List, Dec. 2008, <http://www.top500.org>.
- [62] Systems of Neuromorphic Adaptive Plastic Scalable Electronics (SyNAPSE), Dec. 2008, <http://www.darpa.mil/dso/thrusts/bio/biologically/synapse/>.
- [63] Gluck, K. A., and Pew, R. W. (eds), *Modeling Human Behavior with Integrated Cognitive Architectures*, Lawrence Erlbaum Associates, Mahwah, NJ, 2005.
- [64] Meyer, D. E., and Kieras, D. E., *Precis to a Practical Unified Theory of Cognition and Action: Some lessons from EPIC Computational Models of Human Multiple-Task Performance*, MIT Press, Cambridge, MA, 1986.
- [65] Newell, A., *Unified Theories of Cognition*, Harvard Univ. Press, Cambridge, MA, 1990.
- [66] Avery, E., Kelley, T. D., and Davani, D., "Using Cognitive Architectures to Improve Robot Control: Integrating Production Systems, Semantic Networks, and Sub-symbolic Processing," *Proceedings of the 15th Annual Conference on Behavioral Representation in Modeling and Simulation (BRIMS)*, Baltimore, MD, 2006.
- [67] Kelley, T. D., Avery, E., Long, L. N., and Dimperio, E., "A Hybrid Symbolic and Sub-Symbolic Intelligent System for Mobile Robots," *InfoTech@Aerospace Conference*, Seattle, WA, AIAA, Reston, VA, 2009, AIAA Paper 2009-1976.
- [68] Anderson, J. R., Bothell, D., Byrne, M. D., Douglas, S., Lebiere, C., and Qin, Y., "An Integrated Theory of the Mind," *Psychological Review*, Vol. 111, No. 4, 2004, pp. 1036–1060.
doi: [10.1037/0033-295X.111.4.1036](https://doi.org/10.1037/0033-295X.111.4.1036)
- [69] Anderson, J. R., and Lebiere, C., *The Atomic Components of Thought*, Lawrence Erlbaum Associates, Mahwah, NJ, 1998.
- [70] Hanford, S. D., Janrathitkarn, O., and Long, L. N., "Control of Mobile Robots Using the Soar Cognitive Architecture," *Journal of Aerospace Computing, Information, and Communication*, Vol. 6, No. 6, 2009, pp. 69–91.
- [71] Minsky, M., *The Society of Mind*, Simon and Schuster, New York, NY, 1985.
- [72] Pressman, R., *Software Engineering: A Practitioner's Approach*, McGraw-Hill, New York, NY, 2009.
- [73] Hirai, K., Hirose, M., Haikawa, Y., and Takenaka, T., "The Development of Honda Humanoid Robot," *Proceedings of the IEEE International Conference on Robotics and Automation*, Leuven, Belgium, IEEE, Piscataway, NJ, 1998.
- [74] Kaneko, K., Kanehiro, F., Kajita, S., Hirukawa, H., Kawasaki, T., Hirata, M., Akachi, K., and Isozumi, T., "Humanoid Robot HRP-2," *Proceedings of the IEEE International Conference on Robotics and Automation*, New Orleans, LA, IEEE, Piscataway, NJ, 2004.
- [75] Brooks, R. A., Breazeal, C., Marjanovic, M., Scassellati, B., and Williamson, M. M., "The Cog Project: Building a Humanoid Robot," *Computation for Metaphors, Analogy and Agents*, edited by C. L. Nehaniv, Lecture Notes in Artificial Intelligence 1562, Springer-Verlag, Berlin, 1998.

- [76] Calinon, S., Guenter, F., and Billard, A., “On Learning, Representing, and Generalizing a Task in a Humanoid Robot,” *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. 37, No. 2, 2007, pp. 286–298.
- [77] Hornby, G. S., Takamura, S., Yokono, J., Hanagata, O., Yamamoto, T., and Fujita, M., “Evolving Robust Gaits with AIBO,” *IEEE International Conference on Robotics and Automation*, San Francisco, CA, 24–28 Apr. 2000.
- [78] McCarthy, J., and Hayes, P. J. (eds), *Some Philosophical Problems from the Standpoint of Artificial Intelligence*, Edinburgh Univ. Press, Edinburgh, 1969.

Kelly Cohen
Associate Editor