Presented at the AIAA InfoTech@Aerospace Conference, Seattle, WA, April, 2009
AIAA Paper No. 2009-1949

The Requirements and Possibilities of Creating Conscious Systems

Lyle N. Long[1]
The Pennsylvania State University, University Park, PA 16802

and

Troy D. Kelley[2]
Army Research Laboratory, Aberdeen, MD 20783

This paper discusses the psychological, philosophical and neurological definitions of consciousness and the prospects for the development of a conscious machine in the foreseeable future. Various definitions of consciousness are introduced and discussed within the different fields mentioned. We conclude that a conscious machine 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 architecture 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 we will have systems with intelligence equal to humans by the late 2020's, 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. In this paper we will compare and contrast 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.

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 synonomously, but these are different ideas. 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."

[1] Distinguished Professor, Aerospace Engineering, Bioengineering, and Mathematics, AIAA Fellow.
[2] Engineering Psychologist, Human Research and Engineering Directorate.

A system can be autonomous but not intelligent (e.g. an earthworm) or it could be intelligent but not autonomous (e.g. a supercomputer), as shown in Figure 1. Clearly it is possible to have varying levels of both autonomy and intelligence.

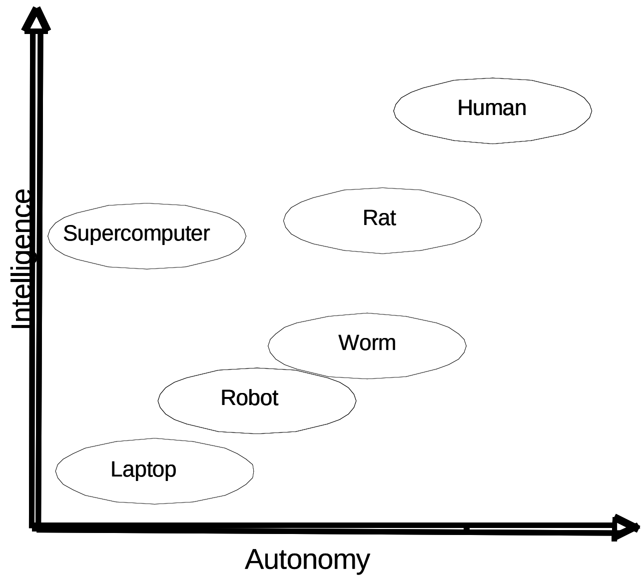

Figure 1. Intelligence vs. Autonomy

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]), electromagnetic fields (as McFadden claims [16]), or quantum effects (as Penrose and others claim [17]). In addition, consciousness is often closely associated with attention [10, 18]. 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 [19]:

> "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. We will not consider esoteric questions such as do we 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 [20].

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 [21]. 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 (and probably consciousness) is show in Figure 2 (from [22]). Current robotic vehicles or systems have probably not achieved the intelligence of a rat or the autonomy of a worm.

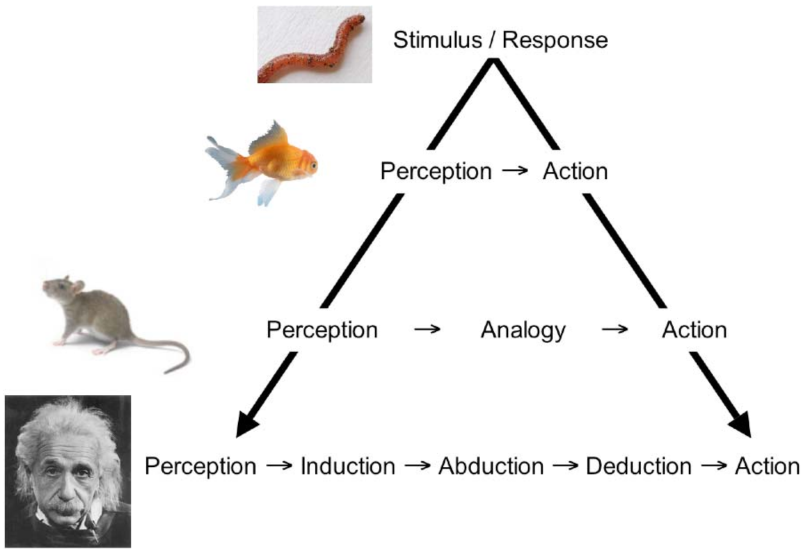

Figure 2. Levels of Intelligence [22]

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. We will briefly review some of the history, key figures, and key publications of consciousness. Most importantly, we will briefly review the current understanding of the 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.) [23] in his discussions of the soul. Descarte (1596-1650) is known for his famous quote "Cogito, ergo sum" ("I think, therefore I am"). Descarte's 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 we know so much about neuroscience. Of course there are people who believe all sorts of things that make no logical sense [24, 25]. The authors of this paper believe that consciousness will be explained completely once we understand more about the brain and nervous system (i.e. materialism and naturalism). Locke (1632-1704) [26], 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 an excellent definition.

More recently there have been several books on consciousness by modern philosophers, including Dennett [4, 6], Searle [27], and Block et al [20]. 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 we believe a conscious machine should be constructed. Dennett also discussed the possibility of conscious machines [6]. 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."

Searle's Chinese Room concept is, in our opinion, an interesting criticism on symbolic systems, and it points out the limitations of symbolic systems which are not bound to experience. Logic or symbolic manipulations alone, however, cannot unravel the mysteries of the brain. 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 [20], and "access consciousness" and "phenomenal consciousness." The "easy" cognitive problem (which isn't necessarily easy at all!) refers to "discriminatory abilities, reportability of mental states, the focus of attention, and the control of behavior." Most philosphers agree that these will eventually be explained via neuroscience. On the other hand, the "hard" problem refers to subjective experiences or "what it is like to be". For example, Nagel's "What is it like to be a bat" [28] . For those interested in building conscious machines, the "hard problem" does not need to be considered. We will also not concern ourselves here with phenomenal consciousness or qualia [20], which Dennett [20], Rosenthal [29] and others argue against.

Seager [30] 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 be eventually built.

In this paper we are content to deal with the more practical and useful "easy problem" and "access consciousness." We also agree that the "easy problem" can be explained by science and those systems can be reverse engineered, and therefore we believe conscious (self-aware) intelligent machines can be built. We will leave the questions about what it feels like to be a cat or to sense the color red to the philosphers.

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."
>
> "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."

Freud [31] said: "The process of something becoming conscious is above all linked with the perceptions, which our sense organs receive from the external world." Gray [19] includes an excellent modern discussion of our current understanding of consciousness from a psychological perspective.

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 [32] discusses the importance of conscious and unconscious processes in describing his global workspace system, which has recently seen some possible experimental support [33].

We also need to differentiate between the process of *developing* consciousness (fetal and human development) and the systems required to *maintain* consciousness.

C. Neuroscience

The most well-known and complete reference to neuroscience is the book by Kandel, Schwartz, and Jessell [11, 12], but the books by Churchland and Sejnowski [11], Koch [10, 34], LeDoux [35] are fascinating works. Crick and Koch [36] discuss the possible neural correlates of consciousness.

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 [35], and Dennett [24]. LeVay [37] 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, 24].

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 don’t 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.

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

- Association pathways (within cortex)
- Commisural 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 [38]. The cortex is a wrinkled sheet on the outer edge of the brain. It is about 0.2 m^2 and about 1 – 4 mm thick. The neocortex has six layers, older parts of the cortex have only 3 layers.

The cortex is just one part of the brain, but it has about 50 – 200 different regions within it, each a separate sub-system. 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 sub-regions):

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

Clearly the brain is a system of systems, and we have begin to understand many parts of it. In the not too distant future, in our opinion, we should be able to reverse engineer 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 [39-41] 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 cortical neurons for several different animals [42]. Humans have roughly 12 billion cortical neurons (5 billion more than chimpanzees), while the entire brain has about 100 billion. Each of these neurons has roughly 30,000 syntactic connections to other neurons, which gives a total of roughly $3 * 10^{14}$ synapses in the human cortex. If you consider the memory potential of a single synapse as 1 byte, then the human cortex has roughly $3 * 10^{14}$ bytes (300 terabytes). A conscious machine will likely need this level of processing power, learning ability, memory capacity, and network interconnections.

Animal	Brain Weight (gm)	Number of Cortical Neurons (10^6)
Human	1350	11,500
Chimpanzee	380	6,200
Cat	25	300
Rat	2	15
Mouse	0.3	4

Table 1. Brain weight and number of cortical neurons.

While the brain is 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 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. In developing conscious machines, it may be most important to consider how does something *become* conscious. Building a conscious system is more like a developing embryo or teaching an infant, which eventually become conscious. We need to consider what functions are critical in the *development* of consciousness.

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 year four, depending on how you define consciousness. Learning is essential since 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 we 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 Artificial Intelligence

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, since 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 eye is a very complex system [43], and more complicated than a typical computer. In reality the human brain is thousands of parallel and inter-connected neural neworks 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 we would say:

> 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/ouput 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 [44] 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 $100M, requires 1000 m^3 of space, weighs 228,000 kg, and requires 3.9 megawatts (while the human brain requires about 0.0014 m^3 of space, weighs 1.3 kg, and needs 20 watts). Figures 3 and 4 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.

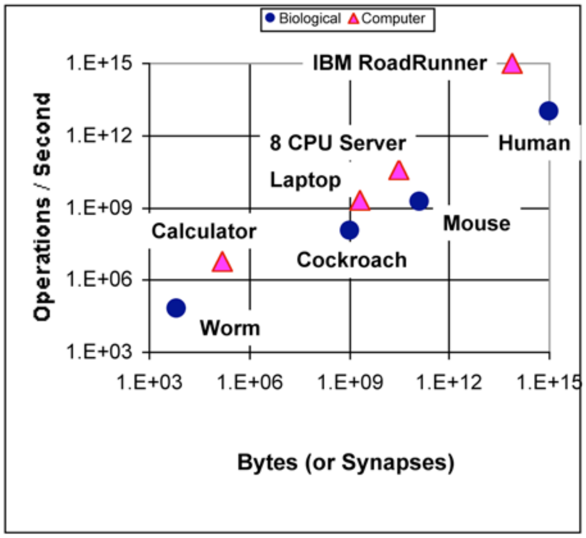

Figure 3. Processing power and memory capabilities of biological and manmade systems.

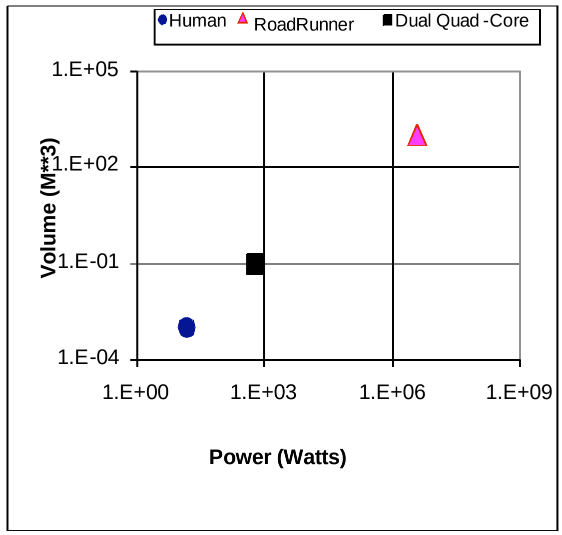

Figure 4. Power and volume requirements for human brain, a server, and a supercomputer.

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 [41], 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 we will need to build hybrid systems. 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 [45] is one such attempt at that.

The field of artificial intelligence (AI) has not produced very intelligent systems, since it has been too focused on symbolic processing. Even cognitive architectures (e.g. [46], Soar [47, 48], SS-RICS [48], and ACT/R [49]), which have been implemented on mobile robots ([50], [48]), are not even close to human intelligence and power and have only rudimentary learning ability. The IEEE Computational Intelligence Society is more focused on sub-symbolic 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 sub-symbolic approaches.

IV. An Engineering Approach to Developing Conscious Systems

Minsky (1988) 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 we *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) allow the successful completion of very complex engineering projects. No single human can fully comprehend a Boeing 777 or the internet, but we can build them. In addition, we know the brain is not made up of that many different parts. There are roughly only 150 different kinds of neurons. We 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, requires 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.

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 sub-symbolic, (e.g. cognitive architectures and neural networks will both be important). We believe it is necessary 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, we feel this might be too labor intensive and it is probably unnecessary for creating a conscious system. Rather, we feel that 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, we believe 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 from human abilities.

The architecture of SS-RICS [48, 51] shown in Figure 5 is one attempt at a general-purpose intelligent system. This system uses a cognitive architecture (based on ACT/R) for decision making but allows for sub-symbolic processing, such as neural networks, for processing sensor data. We have also used the Soar cognitive architecture coupled to sensor inputs and motor outputs [50], as shown in Figure 6. These hybrid approaches will be needed in the near term to build intelligent systems. Figure 7 shows a more general approach and the importance of input data processing.

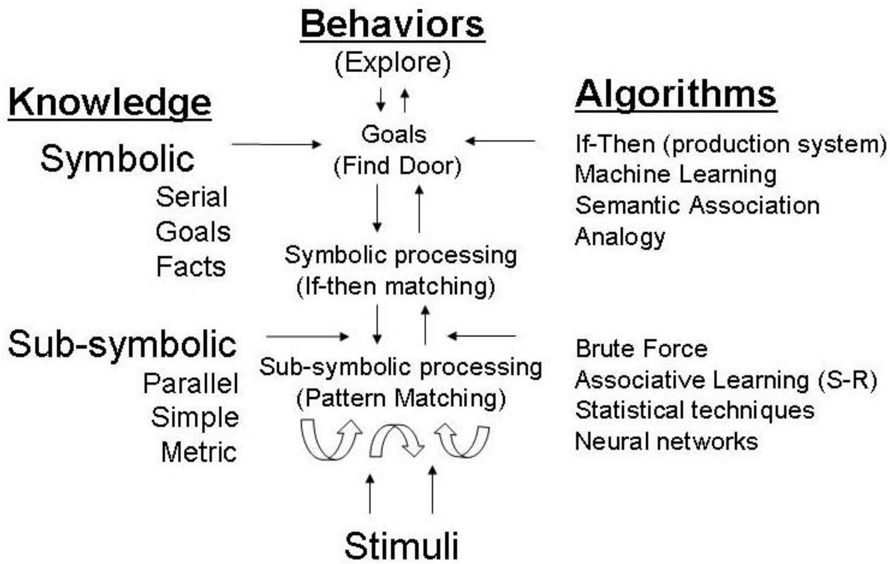

Figure 5. SS-RICS notional data flow emphasizing the symbolic and sub-symbolic distinctions [51]

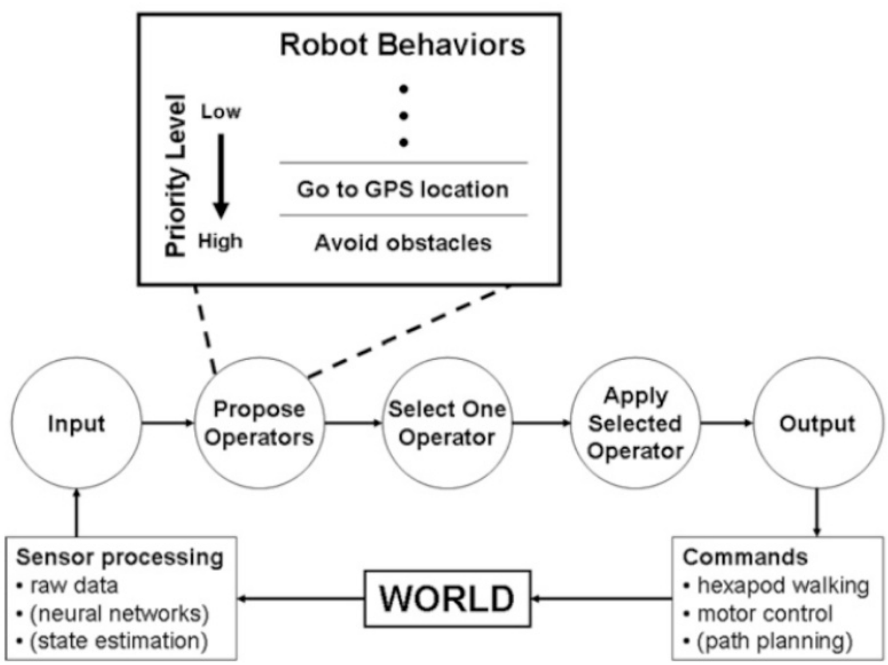

Figure 6. Soar cognitive architecture implemented on a mobile robot [50].

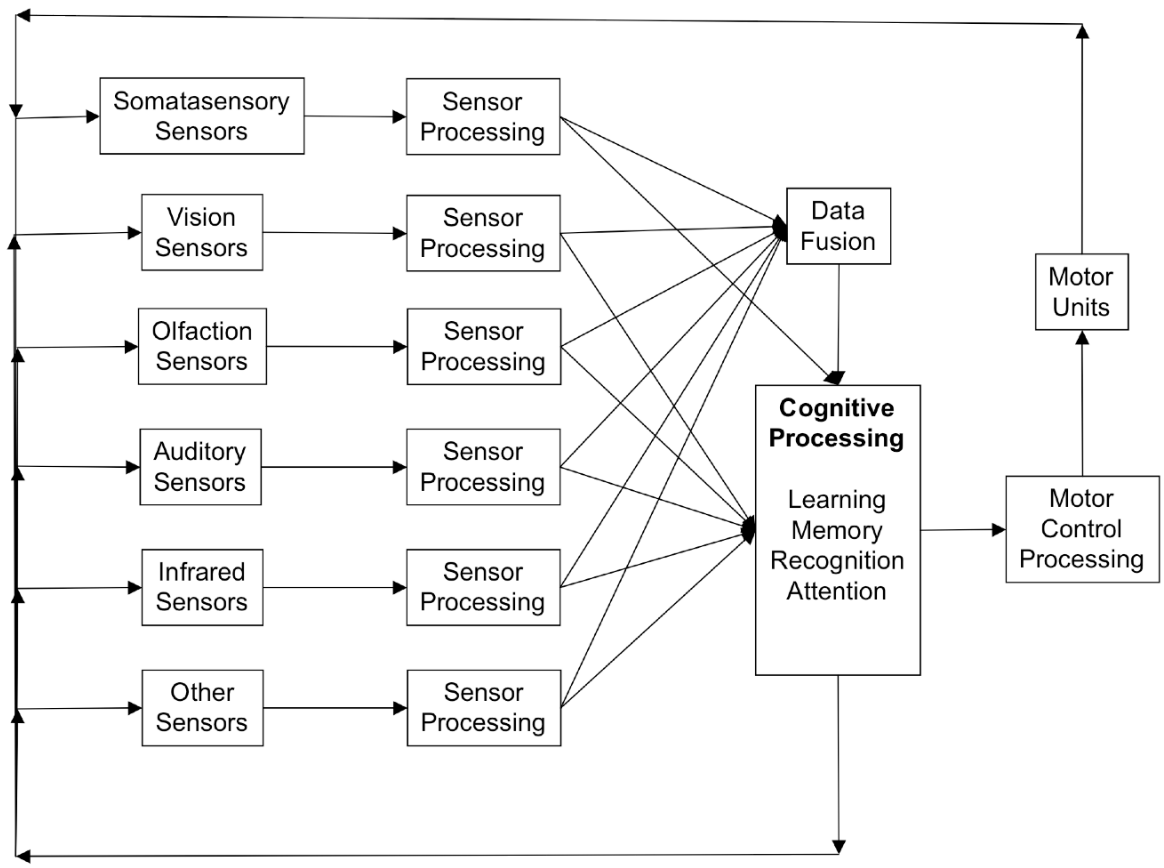

Figure 7. Example architecture for future intelligent (possibly conscious) robots.

Symbolic artificial intelligence (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 require humans to program the rules, and this process is not scalable to billions of rules (a.k.a. the Frame problem [52]). 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.

Conclusion

Conscious machines are likely to be built within this century, if technological advances continue at their current pace, and consciousness can be defined as a computational construct. In order to accomplish this, an interdisciplinary effort involving neuroscience, psychology, computer science, and engineering will be needed. In the near term, it will require a hybrid approach using functional approximations of cognition and will probably include cognitive architectures, neural networks, fuzzy logic, data fusion, parallel computing, and computational power near the levels of the human brain. With proper cognitive development, hardware implementations, and supervised learning, machine consciousness is an achievable goal.

Acknowledgments

Lyle N. Long gratefully acknowledges support as a Moore Distinguished Scholar (2007-2008) at the California Institute of Technology, the Office of Naval Research (Grant No. N00014-05-1-0844), and the Army Research Laboratory (Contract No. TCN 07-305).

References

1. Kurzweil, R. *The Age of Spiritual Machines*: Penguin, 2000.
2. Moravec, H. *Robot: Mere Machine to Transcendent Mind*: Oxford University Press, 1998.
3. Kurzweil, R. *The Singularity is Near*. London: Penquin Group, 2005.
4. Dennett, D. C. *Consciousness Explained*: Back Bay Books, 1992.
5. Descartes, R. *Meditations on First Philosophy*, 1641.
6. Dennett, D. C. "The Practical Requirements for making a Conscious Robot," *Artificial Intelligence and the Mind* Vol. A 349, 1994, pp. 133-146.
7. Pinker, S. *How the Mind Works*: Norton & Company, 1999.
8. James, W. *Principles of Psychology*. New York: Holt, 1890.
9. Edelman, G. M. *The Remembered Present: A Biological Theory of Consciousness*. New York: Basic Books, 1989.
10. Koch, C. *The Quest for Consciousness: A Neurobiological Approach*: Roberts and Company, 2004.
11. Churchland, P. S., and Sejnowski, T. J. *The Computational Brain*: MIT Press, 1992.
12. Kandel, E. R., Schwartz, J. H., and Jessell, T. M. "Principles of Neural Science." 5th ed., 2008.
13. Koch, C., and Tononi, G. "Can Machines be Conscious," *IEEE Spectrum*, 2008.
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.
15. Bekey, G. A. *Autonomous Robots: From Biological Inspiration to Implementation and Control*. Cambridge: MIT Press, 2005.
16. McFadden, J. "Synchronous firing and its influence on the brains electromagnetic field," *Journal of Consciousness Studies* Vol. 9, 2002, pp. 23-50.
17. Penrose, R. *Shadows of the Mind: A Search for the Missing Science of Consciousness*: Oxford, 1996.
18. Koch, C., and Tsuchiya, N. "Attention and Consciousness: two distinct brain processes," *Trends in Cog. Sci.* Vol. 11, No. 1, 2007, pp. 16-22.
19. Gray, P. *Psychology*. New York: Worth Pub., 2007.
20. Block, N., Flanagan, O., and Guzeldere, G. *The Nature of Consciousness: Philosophical Debates*. Cambridge: MIT Press, 1997.
21. Brown, C., Laland, K., and Krause, J., eds. *Fish Cognition and Behavior*. Oxford: Blackwell, 2006.
22. Sowa, J. F., ed. *Categorization in Cognitive Computer Science*: Elsevier, 2006.
23. Caston, V. "Aristotle on Consciousness," *Mind* Vol. 111, No. 444, 2002, pp. 751-815.
24. Dennett, D. C. *Darwin's Dangerous Idea: Evolution and the Meanings of Life*: Simon and Schuster, 1996.
25. Dawkins, R. *The God Delusion*: Houghton Mifflin, 2006.
26. Locke, J. *An Essay Concerning Human Understanding*, 1690.
27. Searle, J. R. *The Mystery of Consciousness*. New York: New York Review, 1990.
28. Nagel, T. "What is it Like to be a Bat?," *Philosophical Review*, 1974, pp. 435-50.
29. Rosenthal, D. M. "How Many Kinds of Consciousness," *Consciousness and Cognition* Vol. 11, 2002, pp. 653-665.
30. Zelazo, P. D., Moscovitch, M., and Thompson, E. *The Cambridge Handbook of Consciousness*. Cambridge: Cambridge University Press, 2007.
31. Freud, S., Strachey, J., and Gay, P. *An outline of psycho-analysis*: Norton, 1989.
32. Baars, B. J. *A Cognitive Theory of Consciousness*. Cambridge: Cambridge Univ. Press, 1993.
33. Gaillard, R., and al, e. "Converging Intracranial Markers of Conscious Access," *PLOS Biology* Vol. 7, No. 3, 2009.

34. Koch, C. *Biophysics of Computation: Information Processing in Single Neurons*: Oxford Press, 1999.
35. LeDoux, J. *Synaptic Self: How Our Brains Become Who We Are*: Penquin, 2003.
36. Crick, F., and Koch, C. "Towards a Neurobiological Theory of Consciousness," *Seminars in The Neurosciences* Vol. 2, 1990, pp. 263-275.
37. LeVay, S. *The Sexual Brain*. Cambridge: MIT Press, 1994.
38. Norden, J. "Understanding the Brain." The Teaching Company, 2007.
39. Long, L. N. "Scalable Biologically Inspired Neural Networks with Spike Time Based Learning," *IEEE Symposium on Learning and Adaptive Behavior in Robotic Systems*. Edinburgh, Scotland, 2008.
40. Long, L. N., and Gupta, A. "Scalable massively parallel artificial neural networks," *Journal of Aerospace Computing, Information, and Communication* Vol. 5, No. 1, 2008.
41. 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.
42. Roth, G., and Dicke, U. "Evolution of the Brain and Intelligence," *Trends in Cog. Sci*. Vol. 9, No. 5, 2005.
43. Martinez-Conde, S., and Macknik, S. L. "Windows on the Mind," *Scientific American* Vol. 297, 2007, pp. 56-63.
44. www.top500.org. "TOP500 Computer List." 2008.
45. http://www.darpa.mil/dso/thrusts/bio/biologically/synapse/. "Systems of Neuromorphic Adaptive Plastic Scalable Electronics (SyNAPSE)." 2008.
46. Gluck, K. A., and Pew, R. W., eds. *Modeling Human Behavior with Integrated Cognitive Architectures*. Mahwah, NJ: Lawrence Erlbaum Assoc., 2005.
47. Newell, A. *Unified Theories of Cognition*. Cambridge: Harvard, 1990.
48. Avery, E., Kelley, T. D., and Davani, D. "Using Cognitive Architectures to Improve Robot Control: Integrating Production Systems, Semantic Networks, and Sub-Symbolic Processing," *15th Annual Conference on Behavioral Representation in Modeling and Simulation (BRIMS)*. Baltimore, MD, 2006.
49. 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.
50. Hanford, S. D., Janrathitikarn, 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.
51. 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*. AIAA Paper No. 2009-1976, Seattle, WA, 2009.
52. McCarthy, J., and Hayes, P. J., eds. *Some Philosophical Problems from the Standpoint of Artificial Intelligence*: Edinburgh Univ. Press, 1969.