2007

Exploring the complex interplay between ai and consciousness

S. Mello

S. Franklin

Exploring the Complex Interplay between AI and Consciousness

Sidney K. D'Mello and Stan Franklin

Department of Computer Science and Institute for Intelligent Systems
209 Dunn Hall, University of Memphis, Memphis, TN 38152, USA
{sdmello|franklin}@memphis.edu

Abstract

This paper embodies the authors' suggestive, hypothetical and sometimes speculative attempts to answer questions related to the interplay between consciousness and AI. We explore the theoretical foundations of consciousness in AI systems. We provide examples that demonstrate the potential utility of incorporating functional consciousness in cognitive AI systems. We also explore the possible contributions to the scientific study of consciousness from insights obtained by building and experimenting with conscious AI systems. Finally, we evaluate the possibility of phenomenally conscious machines.

Introduction

Due to the predominance of the behaviorists in psychology, the scientific study of consciousness was taboo for much of the twentieth century. Thus there was little motivation for AI researchers to include consciousness in such cognitive architectures as SOAR (Laird et al, 1987), ACT-R (Anderson and Lebiere, 1998) or CAPS (Thibadeau et al, 1982). The demise of the behaviorists, together with the roughly concurrent rise of consciousness studies in cognitive psychology and cognitive neuroscience, and the appearance of cognitive architectures such as CLARION (Sun, 1997) and IDA (Franklin et al, 1998) that included consciousness, resulted in the rebirth of the scientific study of consciousness as a subject of possible interest to AI researchers. This new interest gave rise to questions such as:

- Are there theoretical foundations for the study of consciousness in AI systems.
- Can cognitive architectures that include consciousness be of use to AI?
- Can such AI cognitive architectures add to our knowledge of consciousness in humans and animals?
- Are phenomenally conscious AI systems even possible?

This paper embodies the authors' suggestive, hypothetical and sometimes speculative attempts to answer these questions.

Conscious AI Systems

Though consciousness has been the subject of study in philosophy for millennia, it's quite new within AI. What would consciousness even mean in an AI context? The study of consciousness has been divided into the "hard problem" and the "easy problem" (Chalmers, 1996) or, put in another way, the issues of phenomenal consciousness and of access (Block, 1995), or functional (Franklin, 2003) consciousness. Phenomenal, or subjective, consciousness refers to the subjective, individual conscious experience with which we are each so familiar. Functional consciousness refers to consciousness functioning within either natural or artificial systems. Global Workspace Theory (GWT) (Baars 1988, 1997), the currently dominant scientific theory of consciousness, together with its LIDA model (Franklin and Patterson, 2006; Ramamurthy et al 2006) focuses on the function of consciousness in solving the relevance problem, that is, in providing the system with access to its internal resources that are most relevant to the current situation. Hence the alternative name, access consciousness.

There are other theoretical views that draw different distinctions between conscious and unconscious processes as illustrated by Sun & Franklin (2007). The threshold view maintains that the main difference between unconscious and conscious mental representations is that the activation values associated with the former are below a certain threshold (Bowers et al., 1990). A distinct view popularized by the cognitive model ACT-R (Servan-Schreiber & Anderson, 1987) is based on the notion of a chunk – a unitary representation of either a production rule (Rosenbloom et al., 1993) or a sequence of sensory-perceptual elements (Servan-Schreiber & Anderson, 1987). A conscious process would always use several simple chunks while an unconscious process would use a single complex chunk. The dynamical systems view advocated by Mathis and Mozer (1996) claims that the trajectory of conscious processes lands in a stable attractor while unconscious processes are in a transient state.

Saving the question of phenomenally conscious machines for later, we will, for now concentrate on functional consciousness in AI systems. Sun and Franklin have provided a review of AI systems which could conceivably be thought of as embodying functional consciousness (2007), though a number of them don't say so. Another such review will soon appear (Gamez, 2007).

In summary, by a conscious AI system, we will mean an artificial autonomous agent (Franklin and Graesser. 1997) that incorporates some sort of functional consciousness mechanism, that is, an attentional (bring to consciousness) process that allows only the most relevant incoming input through. Every autonomous agent with sufficiently rich sensing in a complex environment will need to filter its input data, often in several different ways. Functional consciousness is one such way as we will see in the next section.

Utility of Conscious Cognitive Agents to AI

As described above functional (or access) consciousness is viewed as the processing of things in experience (Block, 2004) without assuming any subjective experience. Baars (1988, p. 349) highlights several activities that require functional consciousness like prioritization of alternatives, problem solving, decision making, etc. Each of these functional aspects of consciousness can be leveraged to solve critical problems in AI and robotics. For example, Gamez (2007) claims that research devoted to the development of machines with the cognitive characteristics associated with consciousness will ultimately lead to AI systems that are capable of integrating emotions with percepts, attending to various aspects of their environments, and constructing mental representations of virtual scenarios.

In the subsequent discussion we focus on the role of consciousness in solving two problems that are ubiquitous to any artificial agent immersed in a complex, dynamic environment. These include the role of consciousness as a perceptual filter and an important component for human like developmental learning.

Consciousness as a Perceptual Filter

Be it human, animal, software agent or robot, every autonomous agent within a complex, dynamical environment must frequently and cyclically sample (sense) its environment and act on it in order to advance its goals. Due to the inherent diversity and richness of the real world an agent embedded in that world may well require a sequence of finely-tuned filters to sort through the vast amounts of incoming sensory information. The filters would isolate information relevant. to the agents' agenda in order to guide the selection of actions to meet its goals. These filters operate at various levels of granularity and information is filtered through in an increasing order of abstraction in a bottom-up manner.

At the lowest level our sensors have been evolutionarily programmed to act as passive filters for unperceivable signals such as infra red light and ultrasonic sound. One can think of this as a passive filtering process that occurs at the reception level. Next, at the perceptual level an additional phase of content filtration takes place where the sensory input is matched to stored representations that have been experienced in the past and are personally relevant to the agent. Another filtering process also comes into play when one attempts to retrieve information from the vast stores of long term (episodic) memory.

The aforementioned filters operate at the preconscious level. Consciousness plays an important role as a higher-level filter by considering the relevance, importance, urgency, insistence, etc. of information. Consciousness is perhaps the most restrictive filter that agglomerates subconscious, disparate information from sensory processes and memory into a serial, coherent pattern that can be used to recruit resources.

In addition to serving as a perceptual filter, consciousness solves the relevance problem by broadcasting content that is central to the agenda of the agent to the vast array of unconscious processors in order to recruit them into solving the current problem (Baars, 1988).

The benefit of consciousness as a perceptual filter and by providing a solution to the relevance problem has been demonstrated by McCauley in his neural schema mechanism (2002). Simply put, his system performed better with consciousness than without.

Consciousness for Human Like Machine Learning

Dating back to Samuel's checker player (1959), machine learning is among the oldest of the sub-branches of AI with many practitioners and many successes to its credit. Still, after fifty years of effort, there are persistent difficulties that seem to stem from methodological limitations. Machine learning often requires large, accurate training sets, shows little awareness of what's known or not known, integrates new knowledge poorly into old, learns only one task at a time, allows little transfer of learned knowledge to new tasks, and is poor at learning from human teachers.

In contrast, human learning has solved many of these problems, and is typically continual, quick, efficient, accurate, robust, flexible, and effortless. As an example consider perceptual learning, the learning of new objects, categories, relations, etc. Traditional machine learning approaches such as object detection, classification, clustering, etc, are highly susceptible to the problems raised above. However, perceptual learning in humans and animals seem to have no such restrictions. Perceptual learning in humans occurs incrementally so there is no need for a large training set at the front end. Learning and knowledge extraction are achieved simultaneously through a dynamical system that can adapt to changes in the nature of the stimuli perceived in the environment. Additionally, human like learning is based on reinforcement rather than fitting to a dataset or model. Therefore, in addition to learning, humans can also forget. Initially, many associations are made between entities, the ones that are sufficiently reinforced persist, while the ones that aren't decay.

This raises the question of how then do human learn to perceive so robustly and effortlessly? Does consciousness help?

In our view conscious awareness is sufficient for learning. Although subliminal acquisition of information appears to occur, the effect sizes and duration are quite small compared to conscious learning. In a classic study, Standing (1973) showed that 10,000 distinct pictures could be learned with 96% recognition accuracy, after only 5 seconds of conscious exposure to each picture. No intention to learn was needed. Consciously learned educational material has been recalled after 50 years (Bahrick, 1984). No effect sizes nearly as long-lasting as these have been reported in the subliminal learning literature (Elliott & Dolan, 1998). Conscious access greatly facilitates most types of learning.

Once again consider perceptual learning, i.e. the learning of new objects, categories, and relations. In our view humans learn that to which we attend, or what we are conscious of (Baars, 1988, Chapter 9). This occurs by providing the mental (or learned) representation of an object or a category with a measure of its usefulness in the past. The representation is considered to be useful if it is sensed, perceived, and attended to. Hence, while we are profligate in creating new representations, we are saved from computationally intractability by the rapid decay of almost all of the new representations. Only the representations that come to consciousness often and/or at high arousal levels have much chance of not quickly decaying away.

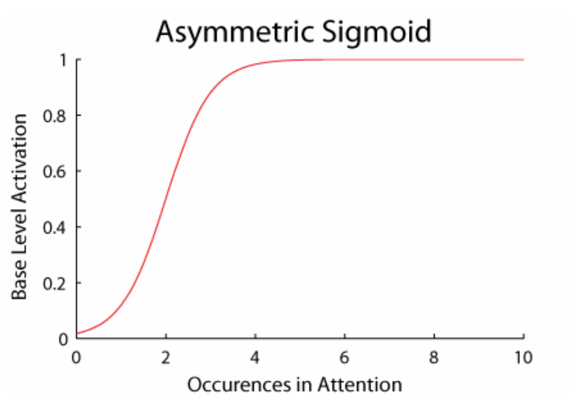

Figure 1. Reinforcement Curve

One way to model the usefulness of a representation is via a saturating, asymmetric sigmoid curve with transition height is a, the transition centre b, and transition width $c\ln(4^{1/d}-1) - c\ln(4^{1/d}3^{-1/d}-1)$ as shown in Figure 1. The X-axis is the number of times the agent is conscious of the stimulus.

As Figure 1 illustrates, the usefulness of the representation (or base level activation) is initially quite small and its initial growth rate is quite slow. However, once a sufficient amount has been accumulated the growth rate is more rapid until it saturates.

AI Can Benefit the Scientific Study of Consciousness

Experimental science progresses via a theorize→predict→experiment→theorize cycle during which a theory is built, predictions made from the theory are tested by experimentation, and the theory is revised in the light of empirical findings, tested again, etc. (Beveridge, 1957, Losee, 1980, Salmon, 1990). All scientific theories are, to some extent, both functional and mechanistic in nature. A *functional* theory describes *what* can be expected to occur in a given situation. A *mechanistic* theory speaks to the *how* of the occurrence, the mechanism that brings it about.

Most current theories of consciousness are typically functional in nature (e.g. Global Workspace Theory, Baars, 1988; 1997; neural synchronization, Crick, 1994; Edelman & Tononi, 2000). These theories are intended to both explain the processes underlying consciousness and predict their functionality, that is, what can be expected to happen under various conditions. Although these functional models are quite useful, even essential, they yield little insight into the *mechanisms* underlying the conscious processes.

In contrast, the control system of any autonomous agent system, by its very nature, must be fully integrated. That is, it must chose its actions based on incoming exogenous or endogenous stimuli utilizing *all* needed internal processes. Also, the control system of an agent must act through its underlying mechanisms. Therefore, the implementation of any autonomous agent requires an underlying mechanistic theory to guide its development.

Almost by definition, the architecture of a functionally conscious agent, an artificial system that employs a cognitive architecture with consciousness to select its next action, is derived from integrated models of the consciousness and cognition of humans and/or other animals. Its control system is designed using that integrated cognitive architecture, and is structurally coupled to its underlying mechanisms. By incorporating consciousness in an autonomous agent, the designer explicitly provides a mechanistic model by fleshing out a functional theory of consciousness.

An AI model builder may also work through something like a theorize→predict→experiment →theorize cycle. A conscious agent is designed and built. It is deployed in an environment and predictions on its behavior are made. Experimentation shows that it doesn't perform as desired. Its functionally conscious control system and, possibly, the underlying mechanisms are redesigned and rebuilt. More experimentation takes place resulting in more redesigning, etc.

The two theorize→experiment →theorize cycles can be amalgamated by means of a conscious agent able to participate in or to replicate a psychological experiment involving human consciousness. The conceptual architecture of the agent would functionally model the conscious process being experimented with on humans or animals. The computational architecture is essentially the same model only acting through the underlying mechanisms. The computational architecture yields insight into the mechanisms underlying consciousness. The human or animal experiments, together with the replicated experiment with the agent, serve to test both the functional model and the associated computational model of consciousness. Both the high-level functional model and the underlying computational model can then be brought more in line with the results of these experiments. After alterations to the agent suggested by the new version of the architecture are made, new psychological experiments can be designed and carried out to test the current version of the theory. The amalgamated cycle continues.

Phenomenal Consciousness in Machines

The question of the possibility of phenomenal consciousness in machines has been debated ad nauseum. Science fiction authors have weighted in on the subject with such characters as Arthur Clark's Hal and Star Trek: The Next Generation's Commander Data. For pointers to the vast literature, please see Blackmore's introduction (2003). For thoughts on the subject from AI researchers and others, please see Holland's edited volume (2003).

One of the current authors (Franklin) is on record doubting phenomenal consciousness in the case of the working software agent IDA who is endowed with functional consciousness (Franklin, 2003). These doubts are based on the lack of convincing evidence to the contrary. Such evidence for other humans consists of verbal reports. In the case of animals, morphological homology in mammals, at least, offers such evidence (Baars, 2001). Further such evidence is offered by binocular rivalry experiments on monkeys (Koch, 2004). For AI systems there is simply no such evidence.

Must we conclude that phenomenally conscious AI systems are simply not possible? By no means. Merker suggests that phenomenal conscious evolved in animals to provide the animal with a stable coherent virtual environment from which to deal effectively with its world (2005). Merker notes that any animal whose sensors move, as do human eyes, must be able to distinguish between actual motion of objects and agents in its environment and their apparent motion caused by the movement of the animal's sensors. According to Merker, one way to solve this problem is by phenomenal consciousness which provides a stable, coherent world to the animal that's independent of the motion of its sensors. He hypothesizes that this solution to the apparent motion problem constituted at least a part of the fitness advantage that led to the evolution of phenomenal consciousness.

Franklin suggests that phenomenal consciousness may provide the same benefits to a cognitive robot, and that the process of producing such a stable, coherent virtual environment in the robot may also result in phenomenal consciousness in a machine (2005). Currently we only know how to produce phenomenal consciousness through biological reproduction. Its mechanisms are only speculated about (Edelman and Tononi, 2000; Dehaene, et al, 2003; Koch, 2004). An attempt to provide a robot with moving sensors with a stable, coherent account of its world may well lead us to explore mechanisms for ignoring apparent motion. These mechanisms, while filtering out apparent motion, may pass through a stable coherent virtual reality that would serve the robot as phenomenal consciousness serves humans. Exploring such mechanisms for filtering out apparent motion seems worth a try.

Discussion

We have demonstrated that there are clear benefits to incorporating consciousness in designing autonomous artificial agents. These include the use of consciousness to provide a filtering mechanism and a potential solution to the relevance problem. Without specifically mentioning consciousness, an earlier AI architecture, the blackboard architecture as used in Hearsay (Hayes-Roth and Lesser, 1977), was noted by Baars (1988, p.79) as a forerunner of Global Workspace Theory and, hence, of functional consciousness. Knowledge sources in the blackboard architecture correspond to processors in GWT. The blackboard's control shell plays the role of attention, the competition for consciousness in GWT. Writing to be blackboard corresponds to the global broadcast. What's missing in the blackboard architecture is effectors capable of external as well as internal actions, and the essential role of consciousness in learning.

Consciousness can also assist in implementing developmental learning mechanisms that are functionally similar to human-like learning. The learning framework that was briefly descried earlier marks a significant departure from current machine learning paradigms. No large training sets would be required. New knowledge would be easily integrated into old. Several tasks could be learned concurrently with transfer of knowledge to new tasks.

It is important to acknowledge that not all artificial agents require consciousness to achieve their goals. In fact agents involved in relatively simple tasks in restricted domains can performs quite well with traditional AI techniques and without any functional consciousness. Expert systems which have been remarkably successful in solving a variety of real world problems serve as exemplars to this point. It is possible that a still more complex artificial autonomous agent with a task requiring more sophisticated decision making would require them. In our view consciousness is best reserved for artificial agents that aspire to be broad models of human cognition and solve real world problems in dynamic environments. It also appears that consciousness comes into its own in agent architectures where online learning of facts and/or skills is of prime importance.

In addition to consciousness being of potential utility to AI, in our view experiments with autonomous artificial agents that incorporate aspects of human consciousness can be leveraged to enhance psychological theories of consciousness. Experiments with conscious agents provide hypotheses that can potentially be verified by conducting experiments with humans. The advantage of hypotheses that stem from AI systems are that they can be specified at sufficient levels of details.

In a similar vein Gamez (2007) suggests the neural correlates of consciousness could be better understood by AI systems that model consciousness with neural networks (e.g. Dehaene et al., 2003; Shanahan, 2006). There are other, more tangible, real world benefits as well. For example, insights gained from such models could also be used to improve the diagnostic capabilities of medical professionals with respect to patients that are comatose or suffering from related phenomena.

We acknowledge that current techniques for studying conscious phenomena at a fine grained level, such as PET, fMRI, EEG, implanted electrodes, etc., are still lacking either in scope, in spatial resolution, or in temporal resolution. PET and fMRI have temporal resolution problems, EEG is well-known to have localizability difficulties, and implanted electrodes (in epileptic patients), while excellent in temporal and spatial resolution, can only sample a small number of neurons; that is they are limited in scope. As a result, many of such emerging hypotheses, while testable in principle, will be difficult to test at the present time. However, improved recording methods are emerging rapidly in cognitive neuroscience. It is important to note that when Global Workspace Theory (Baars 1988, 1997) was first proposed, the core hypothesis of "global activation" or "global broadcasting" was not directly testable in human subjects. Since that time, however, with the advent of brain imaging, widespread brain activation due to conscious, but not unconscious, processes has been found in dozens of studies (see Baars, 2002; Dehaene, 2001). We expect further improvements to make hypotheses provided by experiments with conscious AI systems testable as well.

Conclusion

In summary, we offer tentative conclusions to each of the questions posed in the introduction:

Are there theoretical foundations for the study of consciousness in AI systems? Baars' Global Workspace Theory (Baars, 1988; 1997) and our LIDA conceptual model that fleshes out GTW (Ramamurthy et al, 2006) provide just such a theoretical foundation. We argue this point elsewhere (Franklin et al, in press).

Can cognitive architectures that include consciousness be of use to AI? We claim that the roles of functional consciousness in 1) filtering so as to allow only the most important information through, and in 2) solving the relevance problem by allowing the recruitment of appropriate internal resources, constitute major contributions to AI systems that include it. In addition to IDA (Franklin, Kelemen, and McCauley. 1998; Franklin 2003), both neural schema (McCauley, 2002) and Conscious Tutoring System (Dubois, Nkambou, and Hohmeyer. 2006) provide examples of such AI systems.

Can such AI cognitive architectures add to our knowledge of consciousness in humans and animals? Software agents endowed with functional consciousness can be used to test mechanistic hypotheses derived from their underlying theory by replicating appropriate i*n vivo* experiments. Knowledge about possible underlying cognitive mechanisms should provide guidance to cognitive neuroscientists researching consciousness.

Are phenomenally conscious AI systems even possible? The idea of implementing a stable, coherent virtual world for a cognitive robot, gives some hope that phenomenally conscious cognitive robots might be designed and built.

Acknowledgement

We are grateful to members of the Cognitive Computing Research Group for most useful conversations on the topics of this article. In particular we would like to thank Bernard Baars, Lee McCauley, Aregahagn Negatu, Uma Ramamurthy, and Paul Mouchon.

References

Anderson, J R and Lebiere, C. 1998. *The atomic components of thought*. Mahwah, NJ: Erlbaum.

Baars, B. J. 1988. *A cognitive theory of consciousness*. Cambridge: Cambridge University Press.

Baars, B. J. 1997. *In the theater of consciousness*. Oxford: Oxford University Press.

Baars, B. 2001. Surrounded by consciousness: The scientific evidence for animal consciousness since the mesozoic. In *Consciousness and its Place in Nature: Toward a Science of Consciousness*. Skovde, Sweden.

Baars, B. J. 2002. The conscious access hypothesis: origins and recent evidence. *Trends in Cognitive Science* 6:47–52.

Bahrick, H. P. 1984. Semantic memory content in permastore: fifty years of memory for Spanish learned in school. *Journal Experimental Psychology: General* 113:1-29.

Beveridge, W., 1957. *The Art of Scientific Investigation*, Vintage/Alfred A. Knopf.

Blackmore, S. 2003. *Consciousness: An introduction*. Oxford: Oxford University Press.

Block, N. 1995. On a confusion about a function of consciousness. *Behavioral and Brain Sciences* 18(2): 265–66.

Bowers, K., Regehr, G., Balthazard, C., & Parker, K. 1990. Intuition in the context of discovery. *Cognitive Psychology*, 22*:* 72–110.

Chalmers, D J. 1996. *The conscious mind*. Oxford: Oxford University Press.

Crick, F. 1994. *The Astonishing Hypothesis*. London: Simon & Schuster.

Dehaene, Stanislas, Sergent, Claire and Changeux, Jean-Pierre. 2003. A neuronal network model linking subjective reports and objective physiological data during conscious perception. *Proceedings of the National Academy of Sciences USA*, 100: 8520-8525.

Dubois, D., Nkambou, R. and Hohmeyer, P. 2006. Supporting simulation-based training using a «conscious» tutoring agent. In *Sixth IEEE International Conference on Advanced Learning Technologies (ICALT'06)*:936-938.

Edelman, G. M. and Tononi, G. 2000. *A Universe of Consciousness. How Matter Becomes Imagination*. London: Allen Lane.

Elliott, R., and Dolan, R. J. 1998. Neural response during preference and memory judgments for subliminally presented stimuli: a functional neuroimaging study. *Journal of Neuroscience*. 18:4697-4704.

Engelmore, R., and Morgan, A. eds. 1986. *Blackboard Systems.* Reading, Mass.: Addison-Wesley.

Franklin, S. 2003. Ida: A conscious artifact? *Journal of Consciousness Studies* 10: 47–66.

Franklin, S. 2005. Evolutionary pressures and a stable world for animals and robots: A commentary on Merker. *Consciousness and Cognition* 14: 115–118.

Franklin, S and Graesser, A. C. 1997. Is it an agent, or just a program? A taxonomy for autonomous agents. In *Intelligent agents iii*:21–35. Berlin: Springer Verlag.

Franklin, S, Kelemen, A. and McCauley, L. 1998. Ida: A cognitive agent architecture. In *IEEE conf on systems, man and cybernetics*:2646–2651: IEEE Press.

Franklin, S and Patterson, F. G. 2006. The lida architecture: Adding new modes of learning to an intelligent, autonomous, software agent. In *Idpt-2006 proceedings (integrated design and process technology)*: Society for Design and Process Science.

Franklin, S, Ramamurthy, U., D'Mello, S. K., McCauley, L., Negatu, A. Silva R. L., and Datla, V. in press. LIDA: A computational model of global workspace theory and developmental learning.

Gamez, D. 2007 to appear. Progress in machine consciousness. *Consciousness and Cognition*.

Hayes-Roth, Frederick and Victor R Lesser. 1977. Focus of attention in the hearsay-ii system. In *Proc. 5th int. Jt. Conf artifictal intelltgence*:27–35.

Holland, O. 2003. *Machine consciousness*. Exeter, UK: Imprint Academic.

Koch, C. 2004. *The quest for consciousness: A neurobiological approach*. Englewood, Colorado: Roberts & Co.

Laird, J. E, Newell, A. and Rosenbloom, P. 1987. Soar: An architecture for general intelligence. *Artificial Intelligence* 33: 1–64.

Losee, J, 1980. *A Historical Introduction to the Philosophy of Science*, Oxford University Press, Oxford, UK, 2nd edition.

Mathis D., & Mozer, M. 1996. Conscious and unconscious perception: A computational theory. *Proceedings of the 18th Annual Conference of the Cognitive Science Society*, 324–328.

McCauley, L. 2002. Neural schemas: Toward a comprehensive mechanism of mind: Ph.D. Dissertation, The University of Memphis.

Merker, B. 2005. The liabilities of mobility: A selection pressure for the transition to consciousness in animal evolution. *Consciousness and Cognition* 14: 89–114.

Ramamurthy, U, Baars, B. J, D'Mello S. K., and Franklin, S. 2006. LIDA: A working model of cognition. In *Proceedings of the 7th international conference on cognitive modeling*, ed. Danilo Fum, Fabio Del Missier and Andrea Stocco:244–249. Trieste: Edizioni Goliardiche.

Rosenbloom, P., Laird, J., & Newell, A. 1993. *The SOAR papers: Research on integrated intelligence*. Cambridge, MA: *MIT* Press.

Salmon, W. C., 1990. *Four Decades of Scientific Explanation*, University of Minnesota Press, Minneapolis, MN.

Samuel, A. L. 1959. Some Studies in Machine Learning Using the Game of Checkers. *IBM J. Res. Develop.* 3:210-229.

Servan-Schreiber, E., & Anderson, J. 1987. Learning artificial grammars with competitive chunking. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*, 592– 608.

Shanahan, M.P. 2006. A Cognitive Architecture that Combines Internal Simulation with a Global Workspace. *Consciousness and Cognition,* 15: 433-449.

Standing, L. 1973. Learning 10,000 pictures. *Quarterly Journal of Experimental Psychology* 25:207-222.

Sun, R. 1997. An agent architecture for on-line learning of procedural and declarative knowledge. In *Proceedings of the international conference on neural information processing (iconip'97): Progress in connectionist-based information systems*:766–769. Singapore: Springer Verlag.

Sun, R. and Franklin, S. 2007. Computational models of consciousness: A taxonomy and some examples. In *Cambridge handbook of consciousness*, ed. P D Zelazo and Morris Moscovitch:151–174. New York: Cambridge University Press.

Thibadeau, R., Just, M. A., and Carpenter, P. A. 1982. A model of the time course and content of reading. *Cognitive Science* 6: 157-203.