

Questions Arising from a Proto-Neural Cognitive Architecture

Christian Huyck and Emma Byrne

Middlesex University, UK

c.huyck@mdx.ac.uk

Abstract

A neural cognitive architecture would be an architecture based on simulated neurons, that provided a set of mechanisms for all cognitive behaviour. Moreover, this would be compatible with biological neural behaviour. As a result, such architectures can both form the basis of a fully-fledged AI and help to explain how cognition emerges from a collection of neurons in the human brain. The development of such a neural cognitive architecture is in its infancy, but a proto-architecture in the form of behaving agents entirely based on simulated neurons is described. These agents take natural language commands, view the environment, plan and act. The development of these agents has led to a series of questions that need to be addressed to advance the development of neural cognitive architectures. These questions include long posed ones where progress has been made, such as the binding and symbol grounding problems; issues about biological architectures including neural models and brain topology; issues of emergent behaviour such as short and long-term Cell Assembly dynamics; and issues of learning such as the stability-plasticity dilemma. These questions can act as a road map for the development of neural cognitive architectures and AIs based on them.

Introduction

According to Newell, a unified theory of cognition is a “single set of mechanisms for all of cognitive behavior” (Newell 1990). This will be expressed by a cognitive architecture, the largely stable structure that directs the process of cognition and learning.

In order to derive this unified theory, models of these mechanisms are required. Thanks to the diligent efforts of neuroscientists, there is an understanding of the structure and function of the neuron. From the patient work of psychologists, there is a good understanding of many aspects of human cognition. What is lacking is an understanding of how human cognition arises from the, on average, 130 billion neurons in the human brain, and the 150 trillion synapses in the neocortex alone (Saver 2006). How, for example, does a set of neurons make it possible to locate the source of an itch to a mosquito alighting on one’s foot? How do these or other neurons make it possible to

reach out and swat that same mosquito? The answer to these questions requires a neural cognitive architecture: a largely stable structure that explains the process of cognition and learning in terms of the underlying neural substrate. There is no existing neural cognitive architecture, based on human or simulation-based studies, that follows Newell’s definition.

How, then, to build a model that bridges the gap between the, admittedly partial, understanding of how neurons function and the, admittedly partial, understanding of cognition? The knowledge required to bridge this gap is not easily accessible via existing investigatory techniques in neuroscience. Arrays of implantable electrodes used in primate research can record the spiking behaviour of up to 128 neurons in a living subject (Kipke et al. 2008), thus revealing some of the picture. But these 130 neurons are equivalent to a billionth of the neurons in the human brain. Furthermore, because of the physical constraints of the recording devices, these neurons must be close neighbours. This makes it impossible to model concerted activity among many neurons over even quite small cerebral volumes.

Useful neural cognitive architectures are therefore more likely to arise from *in silico* models that allow the derivation and testing of alternative hypotheses regarding the connection between neurons and behaviour. Such a model would have a two-fold purpose: firstly, it would allow examination of which patterns of neural activity are consistent with which behaviour, at the level of the individual neuron, the level of the network as a whole and all levels of granularity in between. Secondly, it would provide an easily simulated biological unit for the design and development of artificial, intelligent agents.

A better understanding of the neural cognitive architecture of humans will enable the implementation of interesting and eventually full-fledged Artificial Intelligence. This depends on a basic unit of computation, the neural model, but more importantly the units must then be combined so that cognitive behaviour emerges from the simple behaviour of those units interacting with each other and the environment. Unlike biological systems, these simulated systems can be inspected in minute detail with current technology. This inspection supports a method of refinement that can iteratively lead to improved neural cognitive architectures.

This leaves an important question: what should a neural cognitive architecture look like? The authors propose that

a neural cognitive architecture should be able to model the biology of neurons *at least* at the level of the spiking behaviour of individual neurons. There is still much detail that the spiking model obscures. For example, the effects of different neurotransmitters, or of the length and thickness of axonal fibres, is subsumed in the spiking model, and the resulting spiking rates and timings are all that is modelled. However, this abstraction is useful in that no details of the rate, timing or synchronicity of neuron firing are lost, and a degree of simplicity, which allows the simulation of 100,000 neurons on a desktop PC in real time, is gained. It is this approach that the CABot model adopts.

CABot, described below, is one attempt to develop an agent that bridges the neuro-cognitive gap: model neurons are connected in a network of cell assemblies in order that the agent is able to carry out a set of behaviours. The development of agents like CABot, which are based on proto-neural cognitive architectures, is leading toward a better understanding of the neuro-cognitive gap.

CABot aside, some other precursors to neural cognitive architectures exist (e.g. (Rolls 2008; O'Reilly 1996)), but no single set of neural theories explains cognitive behaviour sufficiently to be called a cognitive architecture. Of course, the same can be said of all existing models including popular architectures such as ACT (Anderson and Lebiere 1998); for instance, ACT does not provide a mechanism for sensing, and ACT models that use sensing bolt on some mechanism. These architectures do, however, provide evidence and a means of testing hypotheses about an, as yet to be developed, cognitive architecture.

The precursors to neural cognitive architectures are relatively new compared to symbolic cognitive architectures, and many questions need to be resolved. Below, four categories of questions are explored. Firstly, long posed questions where progress has been made such as the binding and symbol grounding problems. Secondly, biological architecture issues including neural models and brain topology. Thirdly, issues of emergent behaviour, and fourthly, issues of learning.

These questions have arisen during the development of Cell Assembly roBots (CABots). These are agents that function in a virtual environment, take natural language commands, view the environment, and implement plans. The agents are implemented entirely in simulated fatiguing Leaky Integrate and Fire (fLIF) neurons (Huyck 2007).

Aside from a functioning agent that can parse, plan and see, the work on CABot has exposed a set of questions about neural cognitive architectures in general and the CABot model in particular. The answers to these questions will be future steps towards a neural cognitive architecture. This paper presents these questions in the hope that the reader might be inspired to derive some solutions to the problems the authors have encountered. First, however, is a brief description of the CABot agents.

CABots

The central hypothesis of the authors' research is that the best way to develop an AI is to closely follow the human model. This includes developing agents that behave

in a dynamic 3D environment, that correspond to psychological constraints (cognitive architecture), and that correspond to known neural data. Two early versions of the Cell Assembly Robot (CABot1 and CABot2) have already been developed and CABot3 is nearing completion. Code and other information can be found at <http://www.cwa.mdx.ac.uk/CABot/CABot.html>.

All three agents are based in a simple 3D game environment. They view the environment, maintain simple plans, perform actions, and accept natural language commands from the user, who has direct control of their own agent in the environment. CABot1 (Huyck 2008) does this entirely with simulated neurons. Vision follows a known subset of retinal and primary visual cortex behaviour. In all three agents, the neural network is divided into subnetworks; in some cases, like the retina and V1, the subnets correspond to brain areas, but in other cases they are largely for engineering convenience.

CABot2 extends CABot1 by improving the parser and providing a means of learning goal action pairings (Belavkin and Huyck 2008). The parser is a neural cognitive model of parsing (Huyck 2009a). CABot2 does have some minor symbolic hooks, but is almost entirely implemented in fLIF neurons. CABot3 extends CABot2, working in a more complex environment, and integrating texture recognition into the visual system. It also incorporates a simple form of spatial cognitive mapping (Huyck and Nadh 2009). CABot3, like CABot1, is implemented entirely in simulated neurons. A variant of CABot2 is a cognitive model of the probability matching phenomenon in a classical two-choice task (Belavkin and Huyck 2009).

These CABots have limited learning capabilities. The authors' research group plans to focus upcoming efforts on extending and generalizing learning (see the Questions on Learning section).

Newell (Newell 1990) talks of a single unified explanation for all cognitive behaviour. However, there remain several questions in the light of Newell's definition. For example, what might the characteristics of a unified neural cognitive architecture be? What is a viable definition of "all cognitive behaviour"? If a model encompasses the whole of any defensible definition of "all cognitive behaviour" is it then acceptable as a unified neural cognitive architecture? Are there multiple plausible architectures at different levels of granularity?

Whilst the eventual aim of the CABot project is to devise such an architecture, this aim is, as yet, far from being achieved. However, CABot agents are incomplete neural cognitive models that allow the examination of the link between neurons and some types of cognition. In the current case, CABot models the emergence of symbol grounding behaviour from networks of Cell Assemblies built from component fLIF neurons, in response to visual and natural language input.

It is the aim that such partial models based on the CABot framework will eventually be combined to build a complete and unified neural cognitive architecture.

Progress on Long-Standing Questions

The study of neural behaviour and AI is not new. Many questions have been asked before and some progress has been made. Two important issues are variable binding (Fodor and Pylyshyn 1988; Jackendoff 2002) and symbol grounding (Harnad 1990).

The variable binding question involves the use of variables in neural systems; roughly a variable is assigned or reassigned a value. This task is easy for von Neumann architectures where variable instantiation is a primitive. Recent work in this area gives hope that building cognitive architectures based on neural models can be done. Several solutions have been implemented including binding by synchrony (Malsburg 1981), binding by active links (van der Velde and de Kamps 2006), and binding by short-term synaptic change (Huyck 2009b). So, it is clear that binding can be done in a neural system. However, the questions of how and when the brain actually binds are still open.

Similarly, progress has been made on the symbol grounding problem. The problem can roughly be defined as how does the brain learn the meaning of a symbol? Steels points out that the problem has been solved, meaning that symbols have been learned by simulated neural systems (Steels 2007). While it has been shown that symbols can be learned, it remains to be shown all the ways that symbols are learned, or how all symbols are learned. Symbol grounding is important, because there has been a great deal of research on symbolic AI. Unfortunately, these symbolic systems are not well linked to a real world, and tend to be brittle. It is hoped that effectively grounding symbols via experience with a real world will make them much more effective (Kaplan, Weaver, and French 1990), and allow much more effective AI systems.

Indeed, this is a fertile and important area. Humans acquire concepts over a long period of interaction with the environment. A particularly complex set of symbols involve words, and almost all humans manage to learn a language and thousands of words. Other animals also acquire concepts. However, computational systems currently do not do a very good job of concept acquisition. While certain concepts can be learned in limited circumstances, progress has been slow on any particular system grounding a large number of concepts from the environment. The main issue is how can symbols be grounded effectively and generally?

The progress on variable binding and symbol grounding forces a change in the way questions are asked. It is not a question of whether neurons can do these tasks, but becomes a question of how a single system can effectively implement and use symbol grounding and variable binding to behave more effectively; and a question of how biological systems do these tasks.

Questions about Biological Models

When modelling biologically inspired cognition, what is a good basic unit of computation? How does research based on one type of unit inform research based on other units?

The authors argue that a neuron is a good basic unit of computation. There is extensive evidence about how they

perform and it is reasonably easy to get more information about single neuron behaviour. Large populations of model neurons can be simulated in real time.

fLIF neurons are a sound model, and while the authors hope others will use the model, it is neither expected nor hoped that all other researchers will base their work on fLIF models.

A standard comment in the 80s and even today is that neurons are too complex to understand, and research should instead concentrate on non-neural connectionist models (Smolensky 1988). These comments have helped to lead to a range of mechanisms that may help develop understanding of neural processing. However, it is far from clear how these models map to biological reality. For instance, what do Self Organising Maps (Kohonen 1997) say about neural processing? For each non-neural connectionist model, there is at least one question about how this model informs neural cognitive architectures.

Even focusing on models of neurons, there are a range of questions. All models are incomplete, and are limited to a particular level of temporal or spatial granularity of the physical world. Even complex compartmental models of neurons (Hodgkin and Huxley 1952) are imperfect. Moreover some models are so complex that they cannot be simulated in real time for even a single neuron, and a vast number of neurons are needed to perform cognitive behaviours. Consequently, simpler models are important with the Leaky Integrate and Fire (LIF) model (Maas and Bishop 2001) being used by many researchers. Even within LIF models there are a range of models with different time granularities incorporating different behaviours. All of these models tell us something, but there are a range of questions on how information from one model can be used by other models.

The purpose of a model is to offer a simplified view of a complex system. For a successful model, this will be at a level of granularity that is necessarily less complex than the original system, but that is faithful enough that predictions can be made concerning, and generalities extracted from, the original system. With simpler models of neurons that can be efficiently simulated, it becomes possible to examine behaviour emerging from suites of neurons. The fLIF model is sufficiently simple that 100,000 of them can be simulated in real time on a desktop PC, but fine grained enough that questions can be posed about the relationship between assemblies of real neurons and behaviour. The CABot model encompasses many levels of cognition but poses the challenge: given this level of granularity, how can substantially more than 100,000 neurons be simulated in real time?

Another set of questions involves types of neurons. Much simulation work has assumed that all neurons are roughly the same. However, biologists note radically different behaviour between classes of neurons. Which classes of neurons are important, and how do they interact?

Finally, a great deal is known about brain topology, from static tissue analysis and imaging data via, for instance, fMRI, but, this evidence is just suggestive. While connectivity is important, as is brain activity under certain conditions, neither is conclusive. Wiring diagrams are incomplete and difficult to understand; brain areas may be involved in

one task, but also involved in other, seemingly, unrelated tasks; and even understanding of the execution of known tasks is far from complete. Mapping brain areas, laminar architecture and biological neural types to cognitive behaviour leaves a range of unanswered questions.

Questions on Emergent Behaviour

One of the problems facing neuro-biology is that it is difficult with current techniques to effectively measure the fine grained behaviour of a large number of neurons. So it is difficult to see how complex cognitive behaviour emerges from behaving neurons. This is further complicated by the long time scales over which biological neural systems, those of human and other animals, develop. Computational modelling is ideally situated to address the behaviour of a large number of neurons over a long period.

The questions here become much closer to those asked by traditional cognitive architectures. How do particular behaviours emerge from the behaviour of individual neurons? Many emergent algorithms have been proposed leading to a range of exploration. Perhaps the most important is the Cell Assembly.

One long-standing hypothesis about emergent behaviour is the Cell Assembly hypothesis (Hebb 1949), stating that a group of neurons with high mutual synaptic strength is the neural basis of all concepts. This assembly emerges from interaction with the environment via Hebbian learning, and thus accounts for memory acquisition and short-term memory. There is a wide range of biological evidence supporting this hypothesis (e.g. (Abeles et al. 1993)), and much existing research depends on it.

For example, there is research showing how novel categorizers are learned for visual stimuli in simulated neurons arranged like biological brain areas (Knoblauch et al. 2007). While a great deal of research depends on Cell Assemblies (e.g. (Amit 1989)), existing models of them are weak. As assemblies are the neural basis of psychological concepts, they should behave like psychological concepts, yet most models are more or less binary. These memories either exist or do not, are active or are not, are associated or not. On the other hand, psychological long term memories have a wide range of states from formation through loss; short term memories can be more or less active; there are a range of types and strengths of association between concepts; and both short and long term memories have particular behaviours over time. How can model assemblies behave like psychological memories? The authors plan to work on these types of question soon.

It should be noted that most Cell Assemblies in the CABots are orthogonal; that is, they do not share neurons. This has largely been for engineering convenience and other related work has used overlapping assemblies (Huyck 2007). It seems likely that most assemblies in the biological brain are not orthogonal, with neurons instead shared by assemblies both for the purpose of sharing information and increasing capacity. One stream of research that works in highly overlapping assemblies revolves around spin-glass models (Hopfield 1982; Amit 1989), which make use of Statistical Mechanics. Unfortunately, this work is mainly based

on well-connected systems that are biologically unrealistic. A slightly better graph theoretical model (Valiant 2005) uses random connectivity, though random connectivity is also biologically unrealistic. Using small world topologies (Bullmore and Sporns 2009), or direct mappings from biology would be a better approach for initial connectivity. Learning would then be used to set the synaptic weights. The questions that need to be addressed are: What topologies are appropriate and when? How do the initial topologies affect the eventual system?

Cell Assemblies are an ill-defined concept, but they should be able to interact to form more complex structures. For example, humans seem to have rules (with or without variable binding); how do assemblies interact to implement and learn these rules? Humans, and other animals, seem to have more complex structures like cognitive maps; how do they emerge from neurons in collaboration as assemblies?

While Cell Assemblies are long-standing, other emergent algorithms exist. One model of the thalamocortical and corticostriatal system (Granger 2006) shows how a system generates more and more precise categories for an input.

Environmental feedback can be incorporated into learning using neural mechanisms (Belavkin and Huyck 2008). Here interaction with the environment in terms of positive and negative reinforcement can be used to resolve the exploration exploitation dilemma. Separate simulated neural subsystems interact to support this behaviour.

Even work on distribution of knowledge and processing for the frame problem (Shanahan and Baars 2005) is a type of emergent algorithm. Here complex behaviour, resolving the frame problem, is done via simple processors.

There is also an interaction between lower order feature detectors and higher order categorisers. They interact via feedback, but this mechanism is not well understood. How do different sensory brain areas interact to allow the effective development of categorisers?

At this stage, most of these emergent algorithms (Cell Assemblies aside) have very weak biological neural fidelity. This allows exploration of a wide range of possible algorithms, but the lack of biological support makes it more difficult to find the algorithms that are needed.

Questions in this area include: what are the correct emergent algorithms? How do they emerge from neural behaviour? How do they interact with each other?

Questions on Learning

One of the things that has been shown by the CABots is that any process can be implemented with a set of simulated neurons based on a relatively simple neural model. Most of the above questions are a modified version of the question: how do neurons implement a particular behaviour? This overlooks one of the key benefits of neural systems: their ability to learn. Questions about learning abound!

Unlike variable binding and symbol grounding, the authors are unaware of a solution to the long-standing stability-plasticity dilemma (Carpenter and Grossberg 1988). How can a system retain old knowledge yet learn new things? This is linked to psychological behaviour where some mem-

ories are retained and others lost, but some memories (e.g. riding a bicycle) are never lost even if unrehearsed.

However, neurally there seems to be no sound answer. In typical simulations, learning is turned on, then later turned off for testing performance. Even with symbolic architectures like ACT (Anderson and Lebiere 1998) and Soar (Laird, Newell, and Rosenbloom 1987), when learning is left on, there is degraded performance (Kennedy and Trafton 2006). In the brain, neurons are plastic for their entire existence. It is not even clear if synaptic weight changes (i.e. long-term potentiation and depression) are permanent. Perhaps this problem can be addressed by a combination of attractor dynamics, and interaction with the environment. Still no satisfactory solution has been shown to work.

While resolving the stability-plasticity dilemma is complex, a relatively straightforward question is how are synaptic strengths modified? It is widely, but not universally, agreed that this involves Hebbian learning, but within that there are a range of possible rules; which rule or rules are correct? This is an active area of biological research and includes long and short-term potentiation and depression. Still no simple answer is on the horizon and the questions are complicated by the wide range of neural types and presentation mechanisms.

One particularly irksome question involves learning beyond directly stimulated sensory areas. Given Hebbian learning, neurons must co-fire to gain mutual synaptic strength. As most neurons are not directly stimulated by the senses, how can they ever gain synaptic strength from those neurons? That is, how can higher order brain areas ever gain activation in the first place? Spontaneous activation may play a role here (Huyck and Bowles 2004), but the question is not well answered.

Related questions also involve neural and synaptic death and growth. These occur throughout life, but it is not clear how they affect performance.

Learning combined with activation leads to a complex system that has, at least, dual dynamics. This problem is not novel having been raised by Hebb (Hebb 1949). The simplest problem revolves around Cell Assembly formation. Learning (the long dynamic) requires firing (the short dynamic), but a Cell Assembly must fire before it can be learned. The initial firing is caused by environmental stimuli, so the problem is reduced. Still questions about the interactions of the dynamics are still open.

Hebbian rules are generally used to learn new categories, but there are other types of psychological memories. For example, how are episodic memories learned and when are items forgotten? Again the question is not whether a neural system can do this. The question should be how a single system can effectively learn and use episodic memories along with all other necessary mechanisms to behave more effectively? The same question must be asked about how production rules or grammar rules are learned?

There has been a vast amount of research on learning in general, and biologically plausible learning in particular. While progress has and is being made, learning is at the centre of neural cognitive architectures. These questions are of utmost importance.

Conclusion

By following the human neural and cognitive architecture, developing a better understanding of these, and implementing simulations based on these, better AI systems can be developed. The authors' research group has developed a proto-neural cognitive architecture that has been used in the implementation of virtual simulated agents. The development of this architecture and these agents has led to a range of questions. The preceding sections have discussed many of these.

Many of the questions have been relatively broad. For focus, seven questions, arranged from specific to general, are presented.

1. How can more neurons be simulated in real time? Possible answers include neural hardware (Khan et al. 2008) and distributing processing across machines.
2. What properties do fLIF neurons not have that are required for some intelligent behaviour? The fLIF model is relatively simple. Are there aspects that it does not account for that are necessary for any intelligent behaviour? Possible examples include ion capacity and glia behaviour.
3. How do specialised areas interact with other areas? There are specialised subnetworks in the topology of CABot that simplify the analysis of the model. How can this division into subnetworks be used as a way of understanding the way in which real brain areas cooperate and communicate?
4. How can Cell Assemblies behave like psychological memories? The authors hope this can be resolved via slight modifications to the neural model, slight modifications to the learning rule, a better understanding of the dual attractor dynamics, and a larger numbers of neurons.
5. What is a good test for a neural cognitive agent? This is both a specific question (what test next?), and a general one (is the Turing test enough?).
6. How can new knowledge be learned while old knowledge is retained (stability-plasticity)? There are many ways to avoid this dilemma, but solving it will radically further the understanding of neural processing.
7. What are the actual mechanisms to resolve these problems and how do they interact? As stated in the Progress on Long-Standing Questions section, it is important how it is done biologically. This includes how different systems, which have been artificially separated by researchers, interact.

By existing and behaving in an environment, a system is forced to adhere to a broad range of constraints, that biological intelligent agents face. Forcing further neural and cognitive constraints on that agent, places it in a part of a huge (probably infinite) dimensional space that humans inhabit. This is the only portion of that space that is known to exhibit intelligent behaviour (that of humans).

The authors feel that developing neural agents is an excellent way to advance AI, and we encourage others to do so. Even for those not pursuing this method, it is hoped that

the questions raised in this paper will provide some pointers for a road map for the development of these types of neural cognitive architectures, and AIs built around them.

References

- Abeles, M.; Bergman, H.; Margalit, E.; and Vaadia, E. 1993. Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *Journal of Neurophysiology* 70:4:1629–38.
- Amit, D. 1989. *Modelling Brain Function: The world of attractor neural networks*. Cambridge University Press.
- Anderson, J., and Lebiere, C. 1998. *The Atomic Components of Thought*. Lawrence Erlbaum.
- Belavkin, R., and Huyck, C. 2008. The emergence of rules in cell assemblies of flif neurons. In *Proceedings of the Eighteenth European Conference on Artificial Intelligence*.
- Belavkin, R., and Huyck, C. 2009. A model of probability matching in a two-choice task based on stochastic control of learning in neural cell-assemblies. In *Proceedings of the Ninth International Conference on Cognitive Modelling*.
- Bullmore, E., and Sporns, O. 2009. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience* 10:186–198.
- Carpenter, G., and Grossberg, S. 1988. The art of adaptive pattern recognition by a self-organizing neural network. *IEEE Computer* 21:77–88.
- Fodor, J., and Pylyshyn, Z. 1988. Connectionism and cognitive architecture: A critical analysis. *Cognition* 28:3–71.
- Granger, R. 2006. Engines of the brain: The computational instruction set of human cognition. *AI Mag.* 27:2:15–32.
- Harnad, S. 1990. The symbol grounding problem. *Physica D* 42:335–346.
- Hebb, D. O. 1949. *The Organization of Behavior*. J. Wiley & Sons.
- Hodgkin, A., and Huxley, A. 1952. A quantitative description of membrane current and its application to conduction and excitation in nerve. *J. of Physiology* 117:500–544.
- Hopfield, J. 1982. Neural nets and physical systems with emergent collective computational abilities. *Proc. of the Nat. Academy of Sciences USA* 79:2554–8.
- Huyck, C., and Bowles, R. 2004. Spontaneous neural firing in biological and artificial neural systems. *Journal of Cognitive Systems* 6:1:31–40.
- Huyck, C., and Nadh, K. 2009. Multi-associative memory in flif cell assemblies. In *Proceedings of the Ninth International Conference on Cognitive Modelling*.
- Huyck, C. 2007. Creating hierarchical categories using cell assemblies. *Connection Science* 19:1:1–24.
- Huyck, C. 2008. CABot1: a videogame agent implemented in flif neurons. In *IEEE Systems, Man and Cybernetics*.
- Huyck, C. 2009a. A psycholinguistic model of natural language parsing implemented in simulated neurons. *Cognitive Neurodynamics*.
- Huyck, C. 2009b. Variable binding by synaptic weight change. *Connection Science*.
- Jackendoff, R. 2002. *Foundations of Language: Brain, Meaning, Grammar, Evolution*. Oxford University Press.
- Kaplan, S.; Weaver, M.; and French, R. 1990. Active symbols and internal models: Towards a cognitive connectionism. *AI and Society* 4:51–71.
- Kennedy, W., and Trafton, J. 2006. Long-term symbolic learning in Soar and ACT-R. In *Proceedings of the Seventh International Conference on Cognitive Modelling*, 166–171.
- Khan, M.; Lester, D.; Plana, L.; Rast, A.; Painkras, J.; and Furber, S. 2008. SpiNNaker: Mapping neural networks onto a massively-parallel chip multiprocessor. In *Proceeding 2008 International Joint Conference on Neural Networks*, 2849–2856.
- Kipke, D. R.; Shain, W.; Buzsaki, G.; Fetzi, E.; Henderson, J. M.; Hetke, J. F.; and Schalk, G. 2008. Advanced Neurotechnologies for Chronic Neural Interfaces: New Horizons and Clinical Opportunities. *J. Neurosci.* 28(46):11830–11838.
- Knoblauch, A.; Kupper, R.; Gewaltig, M.; Korner, U.; and Korner, E. 2007. A cell assembly based model for the cortical microcircuitry. *Neurocomputing* 70:1838–1842.
- Kohonen, T. 1997. *Self-Organizing Maps*. Springer.
- Laird, J.; Newell, A.; and Rosenbloom, P. 1987. Soar: An architecture for general cognition. *Artificial Intelligence* 33,1.
- Maas, W., and Bishop, C. 2001. *Pulsed Neural Networks*. MIT Press Cambridge, MA.
- Malsburg, C. V. D. 1981. The correlation theory of brain function. Technical report, Dept. of Neurobiology, Max-Planck-Institute for Biophysical Chemistry.
- Newell, A. 1990. *Unified Theories of Cognition*. Harvard University Press.
- O’Reilly, R. 1996. *The Leabra Model of Neural Interactions and Learning in the Neocortex*. Ph.D. Dissertation, Carnegie Mellon University, Pittsburgh, PA.
- Rolls, E. 2008. *Memory, attention, and decision-making: A unifying computational neuroscience approach*. Oxford University Press.
- Saver, J. L. 2006. Time Is Brain—Quantified. *Stroke* 37(1):263–266.
- Shanahan, M., and Baars, B. 2005. Applying global workspace theory to the frame problem. *Cognition* 98:157–176.
- Smolensky, P. 1988. On the proper treatment of connectionism. *Behavioral and Brain Sciences* 11:1:1–22.
- Steels, L. 2007. *The Symbol Grounding Problem has been solved. So what’s next?: Symbols, Embodiment and Meaning*. Oxford University Press.
- Valiant, L. 2005. Memorization and association on a realistic neural model. *Neural Computation* 17:527–555.
- van der Velde, F., and de Kamps, M. 2006. Neural blackboard architectures of combinatorial structures in cognition. *Behavioral and Brain Sciences* 29:1–72.