

Middlesex University Research Repository

An open access repository of

Middlesex University research

<http://eprints.mdx.ac.uk>

Huyck, Christian R. ORCID: <https://orcid.org/0000-0003-4015-3549> (2017) The neural cognitive architecture. 2017 AAAI Fall Symposium Series: Technical Report FS-17. In: AAAI 2017 FALL Symposium Series: Symposium on A Standard Models of the Mind, 09-11 Nov 2017, Arlington, Virginia, USA. ISBN 9781577357940. [Conference or Workshop Item]

Final accepted version (with author's formatting)

This version is available at: <https://eprints.mdx.ac.uk/22926/>

Copyright:

Middlesex University Research Repository makes the University's research available electronically.

Copyright and moral rights to this work are retained by the author and/or other copyright owners unless otherwise stated. The work is supplied on the understanding that any use for commercial gain is strictly forbidden. A copy may be downloaded for personal, non-commercial, research or study without prior permission and without charge.

Works, including theses and research projects, may not be reproduced in any format or medium, or extensive quotations taken from them, or their content changed in any way, without first obtaining permission in writing from the copyright holder(s). They may not be sold or exploited commercially in any format or medium without the prior written permission of the copyright holder(s).

Full bibliographic details must be given when referring to, or quoting from full items including the author's name, the title of the work, publication details where relevant (place, publisher, date), pagination, and for theses or dissertations the awarding institution, the degree type awarded, and the date of the award.

If you believe that any material held in the repository infringes copyright law, please contact the Repository Team at Middlesex University via the following email address:

eprints@mdx.ac.uk

The item will be removed from the repository while any claim is being investigated.

See also repository copyright: re-use policy: <http://eprints.mdx.ac.uk/policies.html#copy>

The Neural Cognitive Architecture

Christian R. Huyck

Middlesex University

c.huyck@mdx.ac.uk

London, UK NW4 4BT

Abstract

The development of a cognitive architecture based on neurons is currently viable. An initial architecture is proposed, and is based around a slow serial system, and a fast parallel system, with additional subsystems for behaviours such as sensing, action and language. Current technology allows us to emulate millions of neurons in real time supporting the development and use of relatively sophisticated systems based on the architecture. While knowledge of biological neural processing and learning rules, and cognitive behaviour is extensive, it is far from complete. This architecture provides a slowly varying neural structure that forms the framework for cognition and learning. It will provide support for exploring biological neural behaviour in functioning animals, and support for the development of artificial systems based on neurons.

Introduction

The mind emerges from the behaviour of the brain, which in turn consists of neurons. This was known when the first cognitive architecture was proposed (Newell 1990), and remains uncontroversial. Due to the simplicity of developing systems of rules, most architectures are based around them. This paper proposes a cognitive architecture based on neurons, a neural cognitive architecture.

It is difficult to program things in simulated neurons but, like rule based systems, neurons are Turing complete (Byrne and Huyck 2010). Consequently, it is possible to specify, design and implement a cognitive architecture with neurons as the foundation. This can then be refined so it will converge to the actual neural cognitive architecture.

It is, of course, a long way from the certainty that minds emerge from neurons to developing working cognitive models, agents, and other high and low level behaviours in simulated and emulated neurons¹. Consequently, the initial architecture should be relatively simple and be developed around two other, slightly more controversial systems (Kahneman 2011): a fast, parallel, implicit, subconscious system, and a slow, serial, explicit, conscious system. In this paper, these will be referred to as the slow system and the fast system, and they will need to be implemented in neurons.

Copyright © 2017, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹Emulated neurons run on neuromorphic hardware, and simulated neurons run on standard hardware.

Development of a neuro-cognitive standard model will not be simple. It takes advantage of human functioning neural and cognitive systems as a guide, and other behaving animals, all of which use neurons for their minds.

Can We Do It Now

Since neurons are Turing complete, it is possible to build a cognitive architecture in neurons. While it is possible to reimplement, for instance, ACT-R (Anderson and Lebiere 1998) in neurons, this reimplementation might not take full advantage of neural processing. None the less, it will be useful to build basic components of the initial architecture in less than biologically plausible ways. Once things are running in neurons, new more biologically plausible systems can be developed.

We Can Develop an Architecture in Neurons

This is a viable technology now. The community knows how to make rule based systems in neurons. Though incomplete, there is a great deal of knowledge about neurons. We know about neural and cognitive learning, and can emulate at least, mouse size brains, in real time now.

One of the benefits of neural processing is that it can be highly parallel. Unfortunately, most modern computers are largely serial. However, neuromorphic hardware is becoming increasingly available. For instance, the SpiNNaker system (Furber et al. 2013) emulates a wide range of neural models in 1ms time steps in real time. One current system consists of 500,000 processors, and can emulate 500 million neurons in real time. This is not enough to emulate a full human brain, but is much larger than the number of neurons in a mouse brain. Models can be developed to run in Nest (Gewaltig and Diesmann 2007) on standard computers though the models with large numbers of neurons will run at less the real time; this allows development of systems on standard serial hardware.

Neurons communicate via spikes. This provides a degree of modularity to facilitate development. In many cases, subsystems can be combined and communicate via spikes. In other cases, subsystems will be combined and share neurons.

Other Neural Architectures

There are already artificial cognitive architectures based on simulated neurons, or on simulated neuron-like units. One

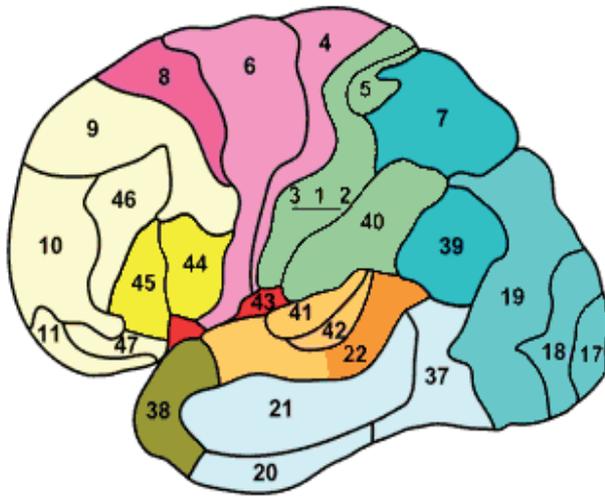


Figure 1: Brodmann Areas, lateral surface of the brain (courtesy of Mark Dubin)

of the best of these is the LEABRA model (O'Reilly 1996). It uses rate coded neurons instead of spike coded neurons, but the two types of systems can be combined. Rate coded neurons do lose information, have problems accounting for synchrony, and have issues with learning.

A second example is the global workspace theory (Shanahan 2006). This uses specialised neural-like units as the basis of its computation. While an interesting exploration of complex behaviours, such as internal simulation, it suffers from its non-biological units.

Another popular neural architecture is the Spaun model (Eliasmith et al. 2012) shows how flexible computation with neurons can be. This benefits from a vector based approach supporting the translation of a range of behaviours into a range of spike based neural models. It suffers from binding via convolution, which does not seem to be biologically plausible, and its dependency on vector representations.

These, and others, are interesting neural systems, and show that it is possible to build a neural cognitive architecture now. However, it is not at all clear which one to build. The brain is self modifying on several time scales, but it is not at all clear how to build a model of even simple point neurons with static synapses. Consequently, we will need to develop a series of prototypes. We do not know everything. We do not know the correct neural models, learning rules, topology, or even brain areas of particular functions.

So, an initial architecture will be based around several subsystems that function in parallel. Like rule based architectures, the subsystems may be customisable. The initial system and its instances will provide direction for the development of future versions.

An Initial Roadmap

This section proposes a relatively simple first version of the architecture. This will consist of customisable neural sub-

systems (descriptions of these subsystems), fully instantiated subsystems (actual neural topologies possibly generated from the descriptions), and fully instantiated full systems (combinations of topologies from the subsystems).

The Slow and Fast Systems

The first subsystem is the slow system. It should be relatively simple to build a relatively unconstrained rule based system in neurons. The key problem is variable binding, and there are several solutions (e.g. (Shastri and Aijjanagadde 1993; van der Velde and de Kamps 2006; Huyck 2009b)). It should be possible to directly translate a particular rule base automatically to neurons. When a particular rule base is converted to neurons, this is a fully instantiated system. Note, that a developer can program the full rule based system directly into neurons, and skip the translation step. (One way to program neurons is to specify the topology using a language like PyNN (Davison et al. 2008).)

While the slow system can be approximated by a rule based system, it is appreciably less clear how the fast system works. It is proposed that it is a form of the spreading activation networks popular in the 80s, e.g. the interactive activation model (Rumelhart and McClelland 1982).

Note that this assumes that a reasonably large amount of the computation the brain does is computation via spreading activation. Bottom up computation, from the senses, is driven by the environment, and is similar to deep nets and other common machine learning connectionist nets. Spreading activation between different sets of neurons is the basis of cognition in the fast system. At the symbolic end of this computation, the spreading activation is determined by the semantics of the symbol, and leads to active symbols (Kaplan, Weaver, and French 1990).

While there has been work in translating general continuously valued neuron models, like interactive activation models, to spiking neurons (Abbott, DePasquale, and Memmesheimer 2016). the author is not aware of a solid theory of how this works in spiking networks. The development of a solid theory of implementing spreading activation nets in spiking neurons is an important challenge for neural computation and cognitive architectures.

The initial fast system will have two largely distinct components: a planning subsystem and an associative memory to represent concepts. Active facts from the associative memory can be used in the planning subsystem to represent the environment, or an internally simulated environment.

One variant of a spreading activation system has been used for planning (Maes 1989) and versions have been built in spiking neurons. In this first component, plans are encoded in the network topology, and are run by spreading activation in response to goals and the environment.

A second component is an associative memory, a type of a semantic net. Concepts and instances of those concepts can be stored in a spreading activation net. Their associations can be stored providing spread of activation behaviour. A variant of this has been used for prepositional phrase attachment resolution, but an extended version is needed.

The fast and slow systems will be combined for the first full system. Facts that support rules will be synaptically

linked to facts from the fast system. So, the systems will run together. A conflict resolution mechanism will determine which action is done when the two systems are in conflict. It is not clear how the brain manages attention (Oberauer and Lin 2017). Inhibitory synapses may focus activation, and synchronous firing may be involved. None the less, some sort of mechanism that is similar to attention will be used to resolve the conflict. One of the key issues of this architecture is attention, its relationship to working memory, and its relationship to the fast and slow systems.

Note that the typical cognitive architecture components of short and long term memory emerge from the neural architecture. Long and short term memory are neural structures (see the Cell Assemblies section). So the initial architecture will involve three general mechanisms, attention, short term and long term memory, emerging from neural behaviour.

Domain Specific Subsystems

Domain specific subsystems will need to be developed. These will need to be integrated with the fast and slow system to form complete initial systems.

Two obvious types of subsystem are sensory subsystems and motor subsystems. It is not difficult to make simple vision systems with spiking neurons. It is however an increasingly difficult task, neurally and non-neurally, to make increasingly sophisticated vision systems.

Other sensory systems include the standard hearing, smell and taste. However, subsystems for balance, internal body sensing (e.g. joint position, and rate of change) will also eventually be needed. In particular, touch is complex, and will need to be integrated with motor action.

Complete systems may need none of these sensory or motion subsystems. For instance, a complete system for conversation about mental arithmetic may only get symbolic input and give symbolic output. None the less, other complete systems will need these sensory systems, for instance a complete simulated driving system would need vision.

Note that to this point in the paper, little has been said about learning. The Learning section will further elaborate, but it is clear that, perhaps, the main reason to use neural systems is that they learn. Even Turing assumed that the first general AI would not be a programmed adult, but a programmed infant that learned. While learning is far from understood, current spiking neural systems are capable of extensive amounts of learning. Moreover, embodied systems that learn an extensive amount about their environments may be industrially viable.

For instance, a spatial cognitive mapping subsystem (Jeffery and Burgess 2006) could be developed. This would make use of grid and place cells. Existing simulated neural systems are able to learn a spatial map of environments, so this would be one form of exploration of learning. This might be extended or varied to make non-spatial cognitive mapping subsystems. Another sample subsystem is natural language understanding. An existing spiking neural subsystem (Huysck 2009a) is already a solid neuro-cognitive model for natural language parsing.

Other subsystems could be developed independently and integrated as time allows. For instance, natural language

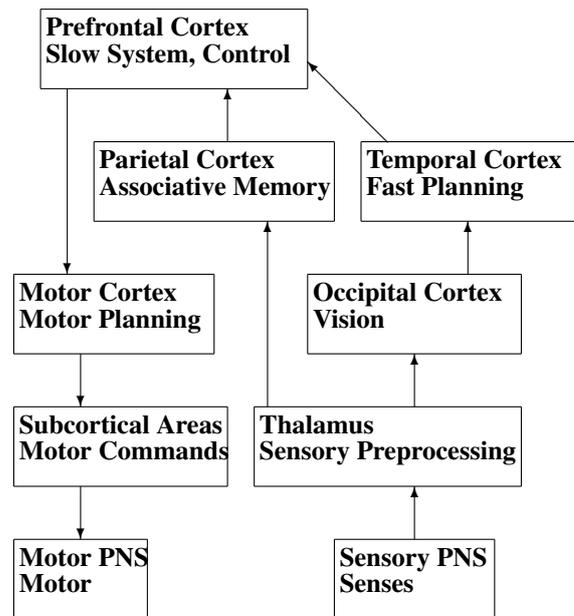


Figure 2: Box Diagram of the Gross Topology of the Architecture. PNS abbreviates Peripheral Nervous System.

generation, arithmetic, episodic memory, and domain specific non spatial cognitive mapping all seem plausible.

Architectural Diagrams

Figure 1 is a representation of the cortex with a standard division. Brodmann areas correspond to physiological features that are constant across most human brains. Work, for example fMRI studies, has been done that suggests some of the cognitive tasks in which each of the areas is involved. It is not entirely clear how these areas correspond to cognitive functionality, so this picture is too specific. Do note that this is just an image of the cortex, while the brain consists of many subcortical levels. It is clear that these subcortical levels are intimately involved in cognition.

A better diagram might be generated from the Human Connectome Project (Essen et al. 2013). The goal of this project is to determine the basic coarse topology of the brain. (The coarse topology refers to the neurons and how they synapse; the fine topology refers to the weights of these synapses.) Another goal is to determine the functional connectivity, showing the effect of these connections. While progress is being made on both of these tasks, unfortunately, the result is far from complete.

Consequently, Figure 2 is a descriptive gross topology diagram showing brain areas, and their hypothesized correlation to the function of the subsystems described above. The arcs refer to the proposed primary activity. Though synapses are unidirectional, if there is a synapse from one area of the brain to another, there are reciprocal synapses (except from the retina). So, the primary activity refers to the hypothesized primary communication of information.

Figure 2 has several major cortical areas. Each box refers to an area of the brain on the first line, and its major function in the architecture in the second line. The prefrontal cortex is the proposed basis for the slow system. Many have proposed that the basis of rule based behaviour is in Basal Ganglia (a portion of the Other Subcortical Areas box). Firstly, it seems unlikely that a slow first system resides entirely in subcortical areas; secondly, subcortical areas connect to all cortical areas, so it is possible that the Basal Ganglia is still involved. Finally, the box diagram refers to the gross topology, so systems will cross boxes. Parietal cortex is responsible for associative memory, and in the initial architecture will be the associative memory component of the fast system. Temporal cortex is responsible for the planning system, occipital cortex is responsible for vision getting input from the retina via the Thalamus. Motor cortex is responsible for motor action, and the Thalamus preprocesses sensory input. Other subcortical areas, beyond the Thalamus, are available for collaboration with all subsystems; early animals do all processing subcortically. Senses and the rest of the peripheral nervous system (PNS) support interaction with the environment.

The overall flow of information is a loop from the sensory systems, through the thalamus to occipital and parietal cortices. Using temporal cortex for support, the prefrontal cortex manages the overall system and passes information to the motor cortex. This passes back through subcortical areas to motor control. This is not the only loop in the system, with the whole system being highly concurrent and recurrent. It is hoped that as the architecture evolves, function will be more closely linked to actual brain topology. Moreover, it is hoped that these draft figures will be replaced by increasingly biologically accurate topologies, at a range of sizes, linked to biological and behavioural data.

Obvious Benefits and Challenges

The previous section provided an early version of the functioning neural architecture. It referred to some of the obvious benefits and challenges of this architecture. The architecture will benefit from work on cell assemblies, embodiment, and existing cognitive architectures. Two of the main challenges are learning and behaviour over a range of time scales.

Cell Assemblies

It has long been hypothesised, that neural cell assemblies are the neural basis of concepts (Hebb 1949), and substantial evidence supports this (e.g. (Harris 2005)). One of the benefits of this mechanism is that the neurons provide both long and short term memories. The long term memory is the formation of the cell assembly. The short term memory is the firing of many of the neurons in the assembly.

While symbol processing systems are powerful, the real strength of animal cognition is the second system. The rapid parallel processing provides the animal with incredible processing. One of the key benefits of the second system is that it provides meaning for the symbols. It grounds them. This is the symbol grounding problem (Harnad 1990).

Cell assemblies provide an intermediate structure consisting of many neurons. Neurons may be in more than one cell

assembly, and assemblies can cross cortical and sub-cortical areas. Depending on the context, particular portions of the assembly may be active or inactive.

Embodiment

A second benefit of the architecture is embodied cognition (Brooks 1991). While complete systems do not need to have complex sensory or motor systems, these systems are readily integrated into the architecture. Vision, touch and other sensory systems can be integrated.

One particular strength is that touch and internal sensory systems can eventually be integrated with the motor systems using learning. The system can, at least theoretically be self modifying. This will allow the robot to improve its performance, and to compensate for changing motor systems. As motors change (as robot developers know they do), the neural control system will modify itself to continue to function.

Also, a great deal is already known about neural sensory and motor behaviour. In particular, it is relatively easy to investigate animal neural behaviour and to use the understanding that has been gained to build neural systems.

Alignment with Existing Architectures

The neural architecture will benefit from work in existing architectures. Existing neural subsystems can be readily integrated within the architecture, communicating with other subsystems via spikes. Below are a few examples of benefits from rule based architectures.

There is a large body of work on the neural implementation of cognitive function. Some of this (e.g. (Jeffery and Burgess 2006)) has been implemented in neurons and this work can readily be integrated with this neural architecture. Other work (e.g. (Koechlin, Ody, and Kouneiher 1997)) merely shows behaviour in the brain, and links this to psychological performance. This can also be implemented in the architecture, but a neural model needs to be developed.

One of the great benefits of ACT-R, is that it has been used to implement many cognitive models. These cognitive models could be re-implemented as neuro-cognitive models. For example, automobile driving behaviour (Salvucci 2006) could be modelled with neurons, reproducing the performance data generated by the original model.

Another issue that has been explored by the ACT-R community is active memory duration. How long does a memory item remain active, how does it decay, and how does it behave if it is reactivated? Duplicating this behaviour in neural memories is an important step.

One of the key concerns with various ACT systems is buffers and buffer sizes. At least initially, this could be tied to neurons in cell assemblies that fire at an elevated rate. Managing amount of neural firing is one of the key problems of neural systems; when too many neurons fire, it is called simulated epilepsy. There are a range of techniques to address this. One is to have four excitatory neurons for every inhibitory neuron, but to give the excitatory neuron four times as much synaptic strength. Another mechanism is to have both local and global inhibitory systems. Local inhibition manages activation in a particular area and helps the system

select a particular cell assembly. Global inhibitory mechanisms take input from large portions of the network, and if there is too much activation, inhibitory neurons fire and reduce the overall activation. Another mechanism is short-term modification of neurons (to prevent them from becoming too active), and short term synaptic modification.

One of the key benefits of Soar (Laird, Newell, and Rosenbloom 1987) is its decomposition of problems into problem spaces. This is readily amenable to decomposition into subsystems. So, there might be a particular neural subsystem for the navigation sub-space, and a different subsystem for analogies.

A second advantage of Soar is chunking. There has been some work in learning rules with neurons, but it is not entirely clear how to do that. Duplicating the Soar mechanism neurally may be a good way to make progress in learning.

Epic (Kieras, Wood, and Meyer 1997) has many solid results in timing of visual task performance. A good visual subsystem, or a good full system with an advanced visual subsystem would duplicate that timing data. Moreover, this data could be directly compared to the firing behaviour of neurons. (Admittedly, this might be ethically difficult to discover in humans, but not in, for instance, rats.)

Learning

There are many types of learning at both the neural and cognitive levels. At the neural level there are different types of synaptic modification including long and short term potentiation and depression, and synaptic growth and death. Similarly, new neurons grow and die throughout life.

From a psychological standpoint, there are also a range of types of learning including semantic learning, episodic learning, different types of sensory learning, motor learning and many others. The memories also vary in the times they are active, and how long they persist.

Learning is one of the key potential benefits of neural systems and one of the key shortcomings of current implementations. Spike Timing Dependent Plasticity (Bi and Poo 1998) is currently in favour as a long term plasticity rule, but even its instantiation is a subject of debate. Ideally the architecture would generate a neural topology that would modify itself at the coarse and fine grained levels, and at various time scales. This evolving topology would be a direct analogy to a living animal whose topology would change differently depending on the environment. Indeed, the changing topology could be directly compared to an animal's.

Unfortunately, the scientific community does not have a clear idea of how to do this. This proposed architecture, however, provides support for exploring these problems. For instance, an episodic memory subsystem could be integrated into a full system. This would inform the development of other memory subsystems, but could also be compared to a second (or third) episodic memory subsystem.

Behaviour of the System Over Many Time Scales

The brain behaves over a range of time scales. Simple neural firing and synaptic communication happens on the ms scale. The neurons and synapses themselves change over short time scales (seconds to minutes) and synapses are modified

(via long-term potentiation and depression) for time scales that persist for months or longer. Sleep is involved in memory formation, but it is not entirely clear how that happens. Some memories seem to persist for years. Also, later memories depend on earlier memories. Initially, only a few of these dynamics will be considered, but eventually, they all need to be considered.

An example of this problem is the stability plasticity dilemma (Carpenter and Grossberg 1988). It must be possible to learn new knowledge, but still retain much of the old knowledge. For instance, humans must be able to learn new categories, while remembering old categories. Things may be forgotten, but important things must be remembered.

Tying Models to Brains and Cognition

One of the key benefits of a neural cognitive architecture over the more traditional variety is that they can be directly compared to brains. While a traditional architecture may say that rules are run in the Basal Ganglia, this architecture can actually compare behaving neurons in the Basal Ganglia to behaving neurons in the model.

Another problem in cognitive neuroscience, and cognitive psychology is development. It is difficult to measure how an animal changes over a period of time, particularly periods like years. Consequently, models are developed that either have a lot of knowledge already in them, or do exceedingly simple tasks. Another benefit of this architecture, is that a full system could persist for years. It is even possible that the system could persist in a virtual environment for virtual years, but in only days of real time.

For the best development of the neural cognitive architecture, it is important that models developed within the architecture, as much as possible, note their biological and neural inaccuracies. For instance, while point neural models (Brette and Gerstner 2005) are models of neurons, they are clearly approximations, and it is not clear what properties a model should have. Similarly, the slow system (proposed in the Initial Roadmap section) may be implemented in neurons, but it is not at all clear that the binding mechanism is correct, and it is almost certain that the neurons will not correspond directly with neurons in any animal.

Finally, this architecture does not require a single topology or even a single type of topology. It is possible to build a range of animal models. So, there may be a mouse brain, a macaque brain and a human brain, and even a crow brain.

Conclusion

Eventually, the actual human and animal neural cognitive architectures will be decoded, providing support for developing AI systems, and understanding both human and animal cognition. Starting the process of understanding the neural cognitive architecture now will speed this understanding.

An initial architecture can be designed around a slow system, implemented via a rule based system; a fast system consisting of associative memory and a planning system, implemented via spreading activation; and other domain specific subsystems. These can all be currently implemented in simulated and emulated neurons.

This initial architecture will be the continuation of a process of discovery of how the brain works. Solutions to large scale problems of neurodynamics, like the stability plasticity dilemma and spreading activation will need to be resolved. Shortcomings in the architecture, and in systems implemented will lead to improvements of this initial architecture. Neuroscientists, neuropsychologists and other scientists can usefully contribute to the architecture, and its connection to underlying neural behaviour.

While it may be possible to build general intelligent systems without using neurons, none currently exist. This architecture will support the development of these systems.

Acknowledgments

This work has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 720270 (the Human Brain Project) and the UK EPSRC. Thanks to David Gamez for comments.

References

- Abbott, L.; DePasquale, B.; and Memmesheimer, R. 2016. Building functional networks of spiking model neurons. *Nature Neuroscience* 19:3:350–355.
- Anderson, J., and Lebiere, C. 1998. *The Atomic Components of Thought*. Lawrence Erlbaum.
- Bi, G., and Poo, M. 1998. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *Journal of Neuroscience* 18:24:10464–10472.
- Brette, R., and Gerstner, W. 2005. Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *J. Neurophysiol.* 94:3637–3642.
- Brooks, R. 1991. Intelligence without representation. *Artificial Intelligence* 47:1:139–159.
- Byrne, E., and Huyck, C. 2010. Processing with cell assemblies. *Neurocomputing* 74:76–83.
- Carpenter, G., and Grossberg, S. 1988. The art of adaptive pattern recognition by a self-organizing neural network. *IEEE Computer* 21:77–88.
- Davison, A.; Brüderle, D.; Eppler, J.; Müller, E.; Pecevski, D.; Perrinet, L.; and Yger, P. 2008. PyNN: a common interface for neuronal network simulators. *Frontiers in neuroinformatics* 2.
- Eliasmith, C.; Stewart, T.; Choo, X.; Bekolay, T.; DeWolf, T.; Tang, Y.; and Rasmussen, D. 2012. A large-scale model of the functioning brain. *Science* 338(6111):1202–1205.
- Essen, D. V.; Smith, S.; Barch, D.; Behrens, T.; Yacoub, E.; Ugurbil, K.; and Consortium, W.-M. H. 2013. The wu-minn human connectome project: an overview. *Neuroimage* 80:62–79.
- Furber, S.; Lester, D.; Plana, L.; Garside, J.; Painkras, E.; Temple, S.; and Brown, A. 2013. Overview of the spinnaker system architecture. *IEEE Transactions on Computers* 62(12):2454–2467.
- Gewaltig, M., and Diesmann, M. 2007. Nest (neural simulation tool). *Scholarpedia* 2(4):1430.
- Harnad, S. 1990. The symbol grounding problem. *Physica D* 42:335–346.
- Harris, K. 2005. Neural signatures of cell assembly organization. *Nature Reviews Neuroscience* 6:399–407.
- Hebb, D. O. 1949. *The Organization of Behavior: A Neuropsychological Theory*. J. Wiley & Sons.
- Huyck, C. 2009a. A psycholinguistic model of natural language parsing implemented in simulated neurons. *Cognitive Neurodynamics* 3(4):316–330.
- Huyck, C. 2009b. Variable binding by synaptic weight change. *Connection Science* 21:4:327–357.
- Jeffery, K., and Burgess, N. 2006. A metric for the cognitive map: found at last? *Trends in Cognitive Sciences* 10:1:1–3.
- Kahneman, D. 2011. *Thinking, fast and slow*. Macmillan.
- Kaplan, S.; Weaver, M.; and French, R. 1990. Active symbols and internal models: Towards a cognitive connectionism. *Connection Science* 4:51–71.
- Kieras, D.; Wood, S.; and Meyer, D. 1997. Predictive engineering models based on the epic architecture for a multimodal high-performance human-computer interaction task. *ACM Transactions on Computer-Human Interaction* 4:3:230–275.
- Koechlin, E.; Ody, C.; and Kouneiher, F. 1997. The architecture of cognitive control in the human prefrontal cortex. *Science* 302:5648:1181–1185.
- Laird, J.; Newell, A.; and Rosenbloom, P. 1987. Soar: An architecture for general cognition. *Artificial Intelligence* 33,1:1–64.
- Maes, P. 1989. How to do the right thing. *Connection Science* 1:3:291–323.
- Newell, A. 1990. *Unified Theories of Cognition*. Harvard University Press.
- Oberauer, K., and Lin, H. 2017. An interference model of visual working memory. *Psych. Review* 124:1:21–94.
- O'Reilly, R. 1996. *The Leabra Model of Neural Interactions and Learning in the Neocortex*. Ph.D. Dissertation, Carnegie Mellon University, Pittsburgh, PA.
- Rumelhart, D., and McClelland, J. 1982. An interactive activation model of context effects in letter perception: Part 2. the contextual enhancement and some tests and extensions of the model. *Psychological Review* 89:1:60–94.
- Salvucci, D. 2006. Modeling driver behavior in a cognitive architecture. *Human Factors* 48:2:362–380.
- Shanahan, M. 2006. A cognitive architecture that combines internal simulation with a global workspace. *Consciousness and Cognition* 15:2:443–449.
- Shastri, L., and Aijanaagadde, V. 1993. From simple associations to systematic reasoning: A connectionist representation of rules, variables, and dynamic bindings using temporal synchrony. *Behaviour and Brain Science* 16:417–494.
- van der Velde, F., and de Kamps, M. 2006. Neural blackboard architectures of combinatorial structures in cognition. *Behavioral and Brain Sciences* 29:1–72.