A Neuropsychological Framework for Advancing Artificial Intelligence

Christian R. Huyck and Hina Ghalib
Middlesex University, UK

Abstract

A framework that combines neural processing, psychological modelling, and Artificial Intelligent (AI) applications is proposed as a means of making steady progress towards the resolution of the problems of the Turing test, understanding neural behaviour, understanding psychological behaviour, and developing useful AI systems. By a principled process of integrating models from these domains, it is hoped that steady progress will be made on a variety of research fronts culminating with a system that can pass the Turing test.

Introduction

Artificial Intelligence (AI) is about building intelligent systems, but unlike humans most existing systems are brittle (Smolensky 1987). The classic test for an AI is the Turing test (Turing 1950), but despite a great deal of effort by the research community and some progress, the community is no where near solving this problem. Consequently, a framework that supports steady progress is highly desirable.

The authors propose that a good way to develop the capacity to build a system that is capable of solving the Turing test is to:

1. Use human-like components (simulated neurons)
2. Keep in touch with reality
   (a) Solve problems like humans do (cognitive architecture)
   (b) Develop interesting systems
3. Integrate learning
4. Take advantage of emergent algorithms
5. Repeat

The basic idea is to focus on the intelligent system that is known, humans and to a lesser extent other animals. Of course, human functioning is not entirely understood, so it is important to work with what is known and in parallel with advances in related fields.

One aspect of human cognitive functioning that is relatively well understood is neurons. Consequently, development should be based on neural models that closely approximate mammalian neural functioning and topology. Similarly, a range of psychological phenomena have been studied, measured, and modelled. Development should generate models and architectures that correspond to known psychological performance, and these models should be based on neurons.

There are a host of neural and psychological behaviours and models. It is relatively simple to build computational models that account for individual behaviours, but what is needed is a system that accounts for them all. Cognitive architectures like ACT (Anderson & Lebiere 1998) account for a large range of behaviours, but are based around symbols. The symbols are not grounded (Harnad 1990), and it is difficult or impossible to learn new symbols that are not just some combination of existing symbols (Frixione, Spinelli, & Gaglio 1989). As neural systems can learn new symbols, a neural cognitive architecture that accounted for a wide range of phenomena would perhaps be more important. Neural cognitive architectures could then be extended beyond symbolic architectures, and could then be used for real world applications.

Indeed developing interesting real world applications is one of the key points of the framework. Systems that interact with complex environments can begin to cope with those environments. Humans are very good at coping with a range of environments, so for a system to be human-like, it must also. Moreover, developing interesting systems allows the construction of systems that are economically useful in the short and medium term. This can sustain the framework beyond the intellectual pursuit.

One key aspect of human, neural, and the best AI systems is that they can learn. The framework thus places learning in a central position. However, at this stage it is not entirely clear how best to build complex neural systems. Cell Assemblies (CAs see the **Neuron** section below) provide one mechanism for an intermediate level of representation, but it is hoped that support may be provided from the architecture to implement emergent algorithms (e.g. (Granger 2006)) that will simplify development of more complex systems.

Initially neural cognitive architectures will be relatively simple. By repeating the development process to extend the architecture and the capabilities, progress can be made in an incremental fashion.

Developing a system that can pass the Turing test in this fashion will also enable the research community to better understand neural functioning, psychological functioning, and will enable the development of AI systems that are incredibly useful. For instance, a system that could pass the Turing test would be an excellent basis for a user interface.

Neurons

The human brain is composed of between 10 billion and one trillion neurons (Churchland & Sejnowski 1999). How these neurons interact to produce cognition is a long standing research problem that is, to some extent, the question on which the authors are focused. Various computational models of neurons have been built and simulated to explain complex cognitive task. There are a host of neural models and non neural connectionist models. While connectionist models (e.g. (Rosenblatt 1962)) may lead to a better understanding of parallel processing in systems that are similar to neural systems, they do not have the advantage of being linked to biological neurons. Consequently, basing a framework on a connectionist system could easily lead to properties that do not occur in human cognition, and simulating human-like cognition is the goal that the author's are pursuing. Neural models on the other hand may need to be revised to more accurately reflect the necessary properties of neurons, but the research community is still finding which properties are necessary and why they are necessary. Choosing the right neural model is important, but progress can be made on a range of models simultaneously. Current neuron models range from biologically faithful compartmental models (Hodgkin & Huxley 1952; Dayan & Abbott 2005; Bower & Beeman 1995) to very simple integrate and fire models (McCulloch & Pitts 1943).

One popular model for brain modelling is the attractor net (Amit 1989). There are a range of attractor nets, but the most commonly used is the Hopfield net (Hopfield 1982). The Hopfield net uses an integrate and fire model like the McCulloch Pitts neurons (McCulloch & Pitts 1943). However, these attractor nets typically use biologically invalid topologies including bidirectional connections and well connected networks. This allows them to take advantage of statistical mechanics, but also means that these systems get stuck in stable states and can not move on.

The model that the authors have been using is a fatiguing Leaky Integrate and Fire model. It is a good model of the biological neuron, but we are pragmatic about it and are still open to changing the model and using different ones. However, the authors have yet to see a reason why this model should not work.

One simplification from biological neurons that this model makes is that it processes at discrete time steps instead of continuous time as in biological networks. This allows a large number of neurons to be simulated and makes the model computationally efficient. Each time step is 10 milliseconds apart. This enables the system to ignore certain biological details like synaptic delay, and the absolute and relative refractory period of neurons. (During the absolute refractory period the neuron does not fire again, no matter how strong the excitatory activation it receives.) The use of discrete time steps may still hide certain useful timing behaviour, and if such behaviours are found, the authors would be happy to switch to a continuous time model.

The large number of neurons in the brain connect with one another through synapses to form complex networks. These synapses are modified with experience and form the basis of learning and memory (see the **Learning** section below). Hebb suggested a synaptic modification mechanism wherein synaptic strength is increased if the pre-synaptic neuron repeatedly and persistently takes part in firing the post-synaptic neuron (Hebb 1949). As a result of this mechanism, groups of correlated neurons which tend to fire together form a strongly connected group of neurons called a cell assembly (CA). A CA once activated remains active, even after the triggering event has been removed, and serves to represent it. CAs, which form the basis of mental representations in the brain, account for neural processes that occur in the brain during thought and learning. There has been growing experimental evidence that such synaptic modification mechanism exist in the brain and form the basis for learning, and memory (Abeles *et al.* 1993; Marr 1969). CAs can serve as the basic processing unit for higher level brain functions and can provide an intermediate level model of cognition ((Huyck 2001) see the **Emergent Algorithms** section below). Hebb's proposed learning mechanism led to the development of various computational models in an attempt to understand how the brain organizes and produces behaviour (Braitenberg 1989; Palm 1990; Sakurai 1998),and CAs have been simulated for decades (e.g. (Rochester *et al.* 1956)).

The fLIF neuron model has also been used for investigation of cell assembly dynamics (Huyck 2008). This model is likely to be used to model a wide range of cognitive processes and gradually lead to the development of a neural cognitive architecture.

Keep in Touch with Reality

The main goal of research in the field of AI is to create systems which behave intelligently like humans. While it is frequently necessary to build small systems to better understand a particular behaviour, it is important to have a mechanism for assuring progress is really being made. Developing the appropriate type of large systems can assure that the paradigm is reasonable.

Perhaps the best way to build systems of human level intelligence is to understand how cognition is achieved by natural systems. Cognitive models are used to simulate human intelligence in a human-like way. Therefore, models of human cognitive processes can help to achieve the goal of AI. As opposed to cognitive models which focus on specific cognitive tasks, cognitive architectures provide a complete theory of human cognition which covers the full range of human cognitive phenomena. Various cognitive architectures exist that provide a framework for building models (Newell 1990; Anderson & Lebiere 1998). These cognitive architectures have been used by researchers to successfully model human behaviour in a wide range of tasks and to develop practical AI systems (e.g. (Wray *et al.* 2005)).

Cognitive Architecture

One way to move forward on developing a system that can pass the Turing test is to build a better cognitive architecture. Unfortunately, most cognitive architectures, including the two most prominent (Newell 1990; Anderson & Lebiere 1998), are symbolic architectures. These have problems with, among other things, learning new symbols that are not merely combinations of existing symbols (Frixione, Spinelli, & Gaglio 1989). It is the hope of the authors that a cognitive architecture based on neurons will resolve this problem and eventually lead to an AI.

The nascent framework has been used to make some initial progress on psychological tasks. An unreported system has been developed to reproduce data on the Stroop task (Stroop 1935; MacLeod 1991). Another version of the system has made progress on the question of the duration of a short-term memory (Passmore & Huyck subm).

The Stroop task is broadly studied with the prototypical example (Stroop 1935) being colour words written in coloured ink. A colour word like blue or red is presented to a subject, and the subject is asked to name this word. If the word is printed in a colour differing from the colour expressed by the word's semantic meaning (e.g. the word *red* printed in *blue* ink), it takes longer to identify the colour of each printed colour name. This is due to difference between the meaning of the word and the perceived colour. This human psychological phenomenon has been implemented in a system based on CAs and fLIF neurons. CAs representing the lexical representation of the word and the print colour, both are sent external stimulation. The CAs representing the underlying semantic meaning of the colours *red* and *blue* are stimulated by the lexical and print colour CAs. When both lexical and print CAs stimulate the same meaning CA it is activated faster than when the stimuli are sent to different meaning CAs. The results that have been simulated roughly agree with human timing data.

A second task involves the duration of an active memory. Short-term memories vary in the time that they remain active. The direct correlate of short-term memory duration in CA theory is how long the reverberating circuit has a large number of neurons firing. A formula exists from ACT-R (Lewis & Vasishth 2005), and a mathematical model of CA dynamics is largely consistent with this (Kaplan, Sontag, & Chown 1991). A series of implementations using fLIF neurons (Passmore & Huyck subm) begins to make progress in modelling these phenomena.

These two systems are examples of neural cognitive models. The Stroop task is solved by developing a particular topology that mimics the data. The memory task is more fundamental, and is about developing a model of CA persistence that emerges from the underlying parameter set and topology. However, the plan is not just to implement simulations that mimic psychological data in fLIF neurons, but to develop a single system that echoes the data for all the tasks, that is, to build a neural cognitive architecture. In this, as yet to be developed, architecture the topology for the Stroop task will be learned, as will the topology for the memory duration task.

It is hoped that this will gradually lead to the development of artificial systems which can produce the full range of human cognitive behaviours. By reproducing core psychological human behaviours, faithful models of larger aspects of human cognition can be developed.

Develop Interesting Systems

Intelligent behaviour evolves due to active interaction with the complex environment (Simon 1969). Therefore, in order to build truly intelligent systems, it is important to establish a link between internal symbols and external real world entities, events and relationships. The problems of assigning symbolic representations of the system to a meaning in the physical world is termed as the symbol grounding problem (Harnad 1990). Symbolic architectures suffer from the symbol grounding problem, and these architectures are as yet unable to learn new symbols and therefore find it difficult to learn new domains. Consequently, systems based on symbolic architectures can be developed for specific domains only. This is a major shortcoming of such cognitive architectures and shows that they would find it difficult to imitate all the processes of human cognition. On the other hand, cognitive architectures based on neural networks provide a potential solution to the symbol grounding problem. A neural cognitive architecture can thus provide an infrastructure for intelligent systems that remain constant across different domains.

As it is important to develop systems that have links to complex environments, a system in a video game environment has been developed (Huyck 2008) using the framework. The system, called CABot1, acts as an agent in the game assisting the user in the game. CABot1 takes input from the environment in the form of a stream of pictures, and input from the user in the form of textual commands. It uses a simulated retina and primary visual cortex to process the picture in a reasonably biologically accurate manner, though the final step of visual processing is a patently inaccurate secondary visual cortex. It processes the text with a context free grammar parser. It uses the textual commands to set goals, and has a subsystem for goal maintenance planning and action. Finally there is an overall control system to that allows the language and action subsystems to take control. All of this is done with just fLIF neurons.

CABot1 does not do much in the way of learning. The current focus of development is learning of actions (Belavkin & Huyck 2008), learning prepositional phrase attachment ambiguity resolution, and a more psycholinguistically accurate natural language parser, all to be integrated into CABot2 in 2008. CABot3, due in May 2009, should include learning of visual items and labelling them along with a start on addressing the symbol grounding problem.

The current plan is to continue with virtual agents for the medium term and developing a conversational agent to further explore linguistically driven actions. Beyond that physical robots are difficult to implement, but it is hoped that they will eventually be implemented using the framework and that this will lead to improved symbol grounding giving the system a better understanding of the semantics of the physical world.

Learning

One of the key aspects of neural networks is that they can learn. Learning has been a keystone in the hopes of developing a system that can pass the Turing test as far back as Turing himself. That is, it is really difficult to build complex systems, so build a system that is complex but one that is complex enough to learn everything else it needs to know.

Learning is of course widely studied both in psychology and in the guise of machine learning, and is an underpinning strength of symbolic cognitive architectures. There are many machine learning algorithms based on connectionist systems, and learning is also studied in biological neural systems (Abeles *et al.* 1993; Schulz & Fitzgibbons 1997).

Developing a neural cognitive architecture that can be used to implement complex agents can make extensive use of learning. Learning enables networks of neurons to build internal models of the environment and perform complex tasks.

The authors have used Hebbian learning mechanism to train networks of fLIF neurons (Huyck 2007). Hebbian learning is an unsupervised learning mechanism in which synaptic modification is based only on the properties and firing behaviour of the pre and post-synaptic neurons. Moreover, it is biologically motivated; most, and perhaps all, learning in biological neurons is Hebbian.

There are a range of proposed Hebbian learning mechanisms, some have stronger links to biological studies than others, but there is no standard agreement on the biological learning rule or learning rules. As the learning problem is both central and poorly understood, this will be an ongoing area of exploration and exploitation within the framework.

Emergent Algorithms

One of the key problems with the framework is how to move from a neural model to a neural architecture. CAs provide some support here as different concepts are supported by sets of neurons, so CAs act as a way of organising the system and thinking about the system (Huyck 2001). This is also supported by a great deal of research on CAs dating back at least to the 1940s.

However, there are two large gaps, one between neurons and CAs, and another between CAs and the full brain. While the number of neurons participating in CAs is debatable, a number around 100,000 is not an unreasonable starting point. The number of neurons in the brain is on the order of 100,000,000,000 and this leaves many orders of magnitude to handle. On the other hand, groups of neurons of size under 100,000 arrange themselves in structures that are not CAs. For example, on-off detectors in the retina are not involved in reverberating circuits.

The gap between neurons and CAs is being explored. There are a range of mechanisms for looking at neural behaviour, and progress is being made. This is made easier by the tractable number of neurons being studied (under 100,000). Still this problem needs to be addressed in a neural cognitive architecture.

The gap between CAs and the full brain is more difficult, perhaps a larger problem, and is certainly less studied. To some extent this is a problem of self-organisation. Putting inputs into a neural system should allow CAs to form via Hebbian learning. However, this only starts to address the question.

Neurons

Synaptic Associations

Neural Collections

Cell Assembly

Cognitive Map

Figure 1: Working Hypothesis on Hierarchy of Structure

Figure 1 provides a working hypothesis on neural structures. It goes from the smallest unit, neuron, to the largest, cognitive maps (Tolman 1948; Jeffery & Burgess 2006). Synaptic associations are centred around a particular synapse and the neurons it connects, but they in turn are affected by other neurons that they are connected to. Neural collections are just different structures than CAs and they are even more ill-defined than CAs. Cognitive maps are collections of CAs, and perhaps neural collections, that allow more sophisticated behaviour. The question at hand is how these can be formed and used. A related question, is are they sufficient to complete the model or are other structures needed?

One way to address this problem is to look at the human brain. Many features are well known. Laminar architecture is prevalant throughout the brain; primary sensory and motor areas are well mapped; areas such as the hippocampus and broca's area are heavily studied; connectivity between brain areas is well mapped. fMRI and other imaging studies on brain areas functioning under different conditions indicate some of the functions of some of these areas. While rapid advancement is being made in understanding the functions of certain parts of the brain, the overall mechanisms are not well understood.

Another approach is to develop models to account for the behaviours and a host of models have been proposed to account for some of the behaviours that might be generated by these areas. Perhaps the best studied are the sensory areas. There is a large amount of information on how these areas function. There are also models of how these areas work (Knoblauch *et al.* 2007). Similarly, there are models for the functioning of other areas, for example the thalamacortical and corticostriatal system is modelled by (Granger 2006).

However it must be noted that brain structures are not like the virtual machines on which modern computation is built (Sloman 1999). The brain has levels of organisation, but the boundaries are fuzzy, and this fuzziness makes the system more flexible. For example, the CA for a given word has neurons in many areas of the brain (Pulvermuller 1999). This enables the CA to influence and be influenced by a large portion of the brain.

The medium term hope is to begin to develop ways of integrating these emergent algorithms to start to develop an architecture. Initially, it is hoped that ways of improving categorisation, signal processing, and memory storage can be integrated to enable large scale storage of memories. A similar goal is to integrate action, action learning, and memory. This area clearly is in need of exploration.

Turn the Wheel

The CABot1 agent (Huyck 2008) is the most sophisticated system that has been developed to date using this framework, and it is far from passing the Turing test. It is however appreciably more sophisticated than the earlier systems that were developed and indicates the progress that has been made. The final key to the framework is to develop new and more sophisticated systems in a directed way.

After some initial work on CAs (Huyck 1999), an initial proposal for the use of CAs as a key point in modelling intelligent systems was made (Huyck 2001). In that paper, several key problems were proposed, and since then work has been done on many of these. For example, CAs are most directly related to categorisation problems, so systems were developed to learn hierarchical categories (Huyck 2007) and to solve real world categorisation problems (Huyck & Orengo 2005). Sequence (Ghalib & Huyck 2007) and rules (Huyck & Belavkin 2006; Belavkin & Huyck 2008) are other complex problems mentioned in (Huyck 2001) that have been addressed by the framework. The rule work addresses the variable binding problem, a core problem for connectionist systems (Malsburg 1986; Jackendoff 2002), and a draft publication extends this work with a novel solution via short-term potentiation.

In some sense these solutions are ways of solving core problems. However, they are theoretical solutions, and need to be scaled up and integrated with each other. Work on CABot1 is a way of combining the components into a larger system, and synergies arise in this system. For instance, once words and semantic items based on vision are formed (e.g. pyramids), it is relatively simple to associate the two to label the semantic item and perform a simple sort of symbol grounding (publication in draft).

This current paper is a revision of the initial framework proposal (Huyck 2001). A new set of problems needs to be identified for candidate solution. The wheel can then be turned to solve these problems.

There are standard connectionist problems, traditional AI problems, and psychology problems. Connectionist problems include the binding problem, the stability plasticity dilemma, and the symbol grounding problem. Traditional AI problems include the frame problem and eventually the Turing test. As mentioned before, cognitive architectures have been extensively used as models for accounting for psychological data, and performance on these and other experiments need to be modelled. In particular, research in the area of attention is needed. Progress has been made on associative memory, but more is needed to handle concepts such as semantic nets and cognitive maps, and to be more accurate psychologically in terms of memory formation, loss, and activation. There is not a shortage of future steps.

Similarly, building systems that solve real world problems, especially if they cover a broad domain, is a great way to move forward. CABot1 is an example of this, but this agent is just a start in the video game domain. Improvements of sensory mechanisms including vision and hearing are needed. Other virtual agents in other domains would be useful. Eventually, physical robots should also be built.

By repeating this process, the underlying neural model can be improved, the cognitive architecture can be developed and refined, and progress is made.

The Next Turn

To a great extent, progress is being made on multiple fronts simultaneously. This of course includes progress within the scientific community as a whole, and this progress needs to be integrated into the developing model. Within the model, progress is made on at least the neural level, the synaptic learning level, the topological level, the emergent level, and the system level. This section contains the start of a proposal that incorporates these levels.

On the neural level, it is hoped that the core neural model will remain for the duration of this next project. It may need to change, or variants may need to be developed and combined, but this not anticipated.

The synaptic learning level occurs at the interface between neurons. It is Hebbian so it is based solely on the behaviours of the pre and post-synaptic neurons. However, this leaves an enormous range of possibilities. The recent use of short-term potentiation in our simulations, has opened a space. In the past, our work on long-term potentiation has required the rule to learn quickly, but this time can now be elongated to allow more stable behaviour.

The topological distinctions that the model has depended on to date are uni-directional and sparse connectivity. Beyond that use of subnets has enabled compartmentalisation for engineering purposes. This compartmentalisation may be continued especially to develop new functionality. However, it is hoped that existing systems will be made more uniform as underlying theory is developed.

At the emergent level, our primary focus will be on developing widely distributed memories. The system has made some use of CAs that share neurons (Huyck 2007), but has been overly reliant on orthogonal CAs. One route may be to follow a graph theoretical approach (Valiant 2005) that supports distribution. From the psychological perspective a related issue is the loss of access to memories, and from the connectionist perspective this relates to the stability plasticity dilemma (Carpenter & Grossberg 1988). This should also incorporate cognitive maps as ways of making CAs work together.

At the system level, the agent will be extended to allow the production of language. This makes it a conversational agent. This step enables a range of complex behaviour including confirmation, and learning via interaction. It will also require better integration of learning and forgetting of associative memory, and better integration of learning rules. In particular, CAs for words will now be used for production and generation. It should also be noted that the Turing test requires a conversational agent, so this system would be a particularly important step toward solving the Turing test.

This proposal is of course only one option for progress. Many small projects exist, but other possible large turns of the wheel include improved vision, speech recognition, modelling attention, and using the system to implement a user interface.

Conclusion

It should be noted that there is a simulation boundary. Even with the simplified fLIF neural model, only a few hundred thousand neurons can be simulated in real time on a standard PC. Advances in processor speed and memory capacity should help, and the system can be distributed. Moreover, advances in neural chip development should lead to a great advancement in speed in the next few years, but capacity is a current problem.

Research in the field of AI aims to build intelligent artificial systems. Advancement in this field has been slower than anticipated; still progress has been made in various areas which can contribute to emulating human intelligence in machines. This paper describes a framework which integrates models of psychological behaviour and information about underlying neural process to develop useful intelligent systems. A framework connecting neural processing and how it produces psychological behaviour can surely help in building artificial system that can learn intelligent behaviour in similar ways to humans. This framework can, in the long run, lead to the design and development of a cognitive architecture, completely built out of neural components.

Look for low hanging fruit. Many researchers look for problems that can be solved relatively easily, but this leaves a range of problems unsolved. The framework described in this paper is building a scaffolding to reach higher hanging fruit. This will enable the solution of a range of more difficult problems. The framework is open and others are encouraged to join, either in discussion or, preferably, in simulation. Hopefully, together we can build a system that can pass the test.

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*.

Bower, J., and Beeman, D. 1995. *The Book of GENESIS.* Springer-Verlag.

Braitenberg, V. 1989. Some arguments for a theory of cell assemblies in the cerebral cortex. In Nadel, Cooper, C., and Harnish., eds., *Neural Connections, Mental Computation*. MIT Press.

Carpenter, G., and Grossberg, S. 1988. The art of adaptive pattern recognition by a self-organizing neural network. *IEEE Computer* 21:77–88.

Churchland, P., and Sejnowski, T. 1999. *The Computational Brain.* MIT Press.

Dayan, P., and Abbott, L. 2005. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems.* MIT Press.

Frixione, M.; Spinelli, G.; and Gaglio, S. 1989. Symbols and subsymbols for representing knowledge: a catalogue raisonne. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, 3–7.

Ghalib, H., and Huyck, C. 2007. A cell assembly model of sequential memory. In *Twentieth International Joint Conference on Neural Networks*, 625–30.

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. *Proceedings of the National Academy of Sciences* 79.

Huyck, C., and Belavkin, R. 2006. Counting with neurons: Rule application with nets of fatiguing leaky integrate and fire neurons. In *Proceedings of the Seventh International Conference on Cognitive Modelling*, 142–147.

Huyck, C., and Orengo, V. 2005. Information retrieval and categorisation using a cell assembly network. *Neural Computing and Applications* 14:282–289.

Huyck, C. 1999. Modelling cell assemblies. Technical report, Middlesex University.

Huyck, C. 2001. Cell assemblies as an intermediate level model of cognition. In Wermter, S.; Austin, J.; and Willshaw, D., eds., *Emerging Neural Architectures based on Neuroscience*. Springer. 383–397.

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*.

Jackendoff, R. 2002. *Foundations of Language: Brain, Meaning, Grammar, Evolution.* Oxford University Press.

Jeffery, K., and Burgess, N. 2006. A metric for the cognitive map: found at last? *Trends in Cognitive Sciences* 10:1:1–3.

Kaplan, S.; Sontag, M.; and Chown, E. 1991. Tracing recurrent activity in cognitive elements(trace): A model of temporal dynamics in a cell assembly. *Connection Science* 3:179–206.

Knoblauch, A.; Kupper, R.; Gewaltig, M.; Korner, U.; and Korner, E. 2007. A cell asssembly based model for the cortical microcircuitry. *Neurocomputing* 70:1838–1842.

Lewis, R., and Vasishth, S. 2005. An activation-based model of sentence processing as skilled memory retrieval. *Cognitive Science* 29:3:375–419.

MacLeod, C. 1991. Half a century of the Stroop effect: An integrative review. *Psychological Bulletin* 109:2:163–203.

Malsburg, C. V. D. 1986. Am I thinking assemblies? In Palm, G., and Aertsen, A., eds., *Brain Theory*. Springer-Verlag. 161–176.

Marr, D. 1969. A theory of cerbellar cortex. *Journal of Physiology* 202:437–470.

McCulloch, W., and Pitts, W. 1943. A logical calculus of ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics* 5:115–133.

Newell, A. 1990. *Unified Theories of Cognition*. Harvard University Press.

Palm, G. 1990. Cell assemblies as a guideline for brain research. *Concepts in Neuroscience* 1:1:133–47.

Passmore, P., and Huyck, C. subm. Models of cell assembly decay. In *IEEE Systems, Man and Cybernetics Society*.

Pulvermuller, F. 1999. Words in the brain's language. *Behavioral and Brain Sciences* 22:253–336.

Rochester, N.; Holland, J.; Haibt, L.; and Duda, W. 1956. Tests on a cell assembly theory of the action of the brain using a large digital computer. *IRE Transaction on Information Theory* IT-2:80–93.

Rosenblatt, F. 1962. *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms*. Spartan Books.

Sakurai, Y. 1998. The search for cell assemblies in the working brain. *Behavioral Brain Research* 91:1–13.

Schulz, P., and Fitzgibbons, J. 1997. Differing mechanisms of expression for short- and long-term potentiation. *Journal of Neurophysiology* 78:321–334.

Simon, H. 1969. *The Sciences of the Artificial*. MIT Press.

Sloman, A. 1999. Architecture-based cocneptions of mind. In *11th International Congres of Logic, Methodolgy and Philosophy of Science*.

Smolensky, P. 1987. Connectionist AI, symbolic AI, and the brain. *Artificial Intelligence Review* 1:95–109.

Stroop, J. 1935. Studies of inteference in serial verbal reactions. *Journal of Experimental Psychology* 18:643–662.

Tolman, E. 1948. Cognitive maps in rats and men. *Psychological Review* 40:40–60.

Turing, A. 1950. Computing machinery & intelligence. *Mind* 59:236:433–60.

Valiant, L. 2005. Memorization and association on a realistic neural model. *Neural Computation* 17:527–555.

Wray, R.; Laird, J.; Nuxoll, A.; Stokes, D.; and Kerfoot, A. 2005. Synthetic adversarites for urban combat training. *AI Magazine* 26:3:82–93.