Two Simple NeuroCognitive Associative Memory Models

Christian R. Huyck (c.huyck@mdx.ac.uk)
Department of Computer Science
Middlesex University
London UK

Yuhue Ji (YJ097@live.mdx.ac.uk)
Department of Computer Science
Middlesex University
London UK

Abstract

Human memory is associative and emerges from the behaviour of neurons. Two models, based on commonly used biological neural models are presented. The first model uses static synapses to approximate timing behaviour for a Stroop task with congruent conditions responding faster than incongruent conditions. The second model uses plastic synapses to learn a semantic net; it then duplicates the behaviour of a question answering task. This behaviour not only answers correctly, its times are similar to that of human subjects. These models are flawed in many ways, for instance, they use hundreds of neurons instead of the billions of neurons in the brain. They are thus not proposed as anything near a complete final model, but instead as early steps toward the development of more sophisticated neurocognitive associative memory models.

Keywords: Associative Memory; Cell Assembly; Hebbian Learning; Spiking Neuron.

Introduction

Human memory is associative in its nature (Anderson & Bower, 1973). Concepts are associated with related concepts, with, for instance, *Dog* associated with *Bone* and *Canine*. Similarly, human cognition is a product of the behaviour of the brain in general, and neurons in particular. Since associative memory is a key component of the human mind, and the mind is a product of the activity of neurons, developing simulations of associative memory in neurons is important.

One modern description of associative memory is the semantic net (Quillian, 1967). This symbolic representation has been widely used in knowledge representation schemes in Artificial Intelligence. Sub-symbolically, many connectionist and simulated neural systems have been developed to account for associative memory (for a review see (Lansner, 2009)). The authors are, however, unaware of any cognitive models of associative memory based on spiking neurons. As associative memory emerges from the behaviour of spiking neurons, such a model is important.

What is a good way to evaluate a cognitive model of associative memory? This paper describes two simple neurocognitive models of associative memory. The first accounts for the Stroop effect (Stroop, 1935) using static synapses. The second accounts for some subcategorisation hierarchy effects that subjects (Collins & Quillian, 1969) show in answering questions; this model makes use of plastic synapses, but has a rigid training regime.

These neuro-cognitive models are simple. They are based on simple neural models, and the second uses a simple Hebbian learning rule. As they are simple, they are flawed. These flaws are discussed, along with relatively simple mechanisms to add to build better neuro-cognitive models of associative memories. The conclusion includes future work in this area.

Literature Review

The task modelled in this paper is associative memory. There has been a great deal of associative memory modelling and psychological exploration of associative memory, and there is evidence that a crucial component of this memory at the neural level is the Cell Assembly (CA). Below this psychological and neuropsychological area is reviewed along with the neurobiological area of synaptic modification, and neural modelling.

Associative Memory

Human memory is associative. Concepts do not exist in isolation, but in a network of associations. A semantic net is a symbolic representation of this memory (Quillian, 1967). Collins and Quillian ran a psychological study that supports this (Collins & Quillian, 1969). In this study, subjects were given a statement such as *a canary is yellow*, and asked to say if it was true or false. It took subjects longer to respond false, but they also took longer if the fact was associated with a super-category of the item. They hypothesised that *bird* was the super-category of *canary* and *animal* the super-category of *bird*. It took longer to respond to *a canary has skin* than to *a canary can fly*, and the shortest was *a canary is yellow*. Similarly, the time to answer a question about direct hierarchical relations were longer the higher up the hierarchy. Shortest was *a canary is a canary*, followed by *a canary is a bird*, followed by *a canary is an animal*.

The Stroop effect (Stroop, 1935) is a consequence of associative interference. The original task has a colour name presented in a coloured ink. So, the word *blue* might be presented in *red* ink. One task is for the subjects to name the colour of the ink, in the example *red*. If the word and ink are congruent, the subjects respond faster and make fewer errors than if they are incongruent. This is a well known and a well studied phenomenon (MacLeod, 1991). Associative interference applies to many domains beyond colours and is a

window into human associative memory.

Cell Assemblies

The CA hypothesis is that the CA is the neural basis of, among other things, concepts (Hebb, 1949). A CA is a group of neurons that has relatively high synaptic connectivity, and relatively highly weighted synaptic connectivity. Thus, once the CA starts to fire, there is a cascade of firing that causes many of the other neurons in the CA to fire. This firing is the neural basis of a psychological short-term memory. The synaptic change required to make this connectivity is a long-term memory. Hebbian learning naturally leads to this type of structure.

There is a large community of researchers that, in essence, assumes that the CA hypothesis is correct, and the authors include themselves in this category. Though Hebb merely theorised the CA based on the limited biological evidence to hand, significant neurobiological evidence for CAs (see (Huyck & Passmore, 2013) for a review) has accumulated in subsequent decades. Indeed the authors are unaware of any evidence contradicting the CA hypothesis. However, the brain's complexity leaves the exact nature of CAs unclear.

Simulated spiking neurons are powerful computational devices. It is relatively simple to build systems based on spiking neurons that are incompatible with the CA hypothesis. These systems are suspect as models of human psychological behaviour.

Synaptic Plasticity

In computational neuro-biological circles, the most popular learning rule is currently spike timing dependent plasticity (Bi & Poo, 1998) (STDP). It is Hebbian; that is, the synaptic weight is increased if the pre-synaptic neuron tends to contribute to the post-synaptic neuron firing. There are many approaches to developing computational models of STDP.

There are also a wide range of learning mechanisms beyond STDP. One particularly useful system (Zenke, Agnes, & Gerstner, 2015) uses several learning rules to learn stable CAs. The network consists of both inhibitory and excitatory spiking neurons. In their simulations, STDP alone leads to unstable systems and stored CAs are erased over time. There is a rule that depresses synapses at high firing rates, and a related rule that increases synapses at low firing rates; both are based solely on one neuron, and are thus non-Hebbian. There are short term potentiation and depression rules. Metaplasticity rules are explored, and there are a range of time dynamics.

This enables the system to not only retain stored CAs, but to learn new CAs. This addresses the long standing neural stability plasticity dilemma (Carpenter & Grossberg, 1988). If learning remains on, so that new things can be learned, the old memories, stored synaptically, can be erased.

The range of synaptic weight modification mechanisms shows biology contrasts and complements mathematics. Simple mathematical rules help to explain the mechanisms, and can be implemented readily. They are usually approximations to biological mechanisms that are still poorly understood. In particular, the dynamic nature of the neural system, with spiking effecting synapse weights and synapse weights influencing spiking, makes it difficult to understand.

One simple rule is known as Oja's rule (Oja, 1982). This rule leads to the synaptic weight reflecting the likelihood that the post-synaptic neuron fires when the pre-synaptic neuron fires, their correlation. It has two components, an increase rule when the neurons co-fire, and a decrease rule when the pre-synaptic neuron fires, but the post-synaptic neuron does not. It can be modelled as the early part of equations 1 and 2.

$$\Delta_{+} w_{ij} = R * [(1 - w_{ij}) * 2^{(W_B - W_k)}] \quad (1)$$

$$\Delta_{-} w_{ij} = -R * [w_{ij} * 2^{(W_k - W_B)}] \quad (2)$$

In these equations R is the learning rate, and w_{ij} is the current synaptic weight. The exponential components of the equations are the compensatory modifiers, not used in Oja's rule and explained below. Using these rules, if the post-synaptic neurons fires 40% of the times when the pre-synaptic neuron fires, the weight will be approximately 0.4.

In addition to Oja's correlation component, the equations have a compensatory component that forces the total synaptic weight of a neuron toward a value W_B. The compensatory modifier is the exponential value at the end of the equations, and current synaptic weight of a neuron is W_k. The simulation below uses a pre-compensatory rule, where only the total synaptic weight of the pre-synaptic neuron is considered in the weight update.

The compensatory rule initially speeds learning, but also limits the synaptic weight. This limit prevents runaway synaptic growth. The authors have used these rules extensively (Huyck & Mitchell, 2014), and it has been suggested that compensatory processes are required for Hebbian learning (Zenke & Gerstner, 2017).

Simulating Neurons

There are many different computational models of neurons (see (Brette et al., 2007) for a review). One widely used class of model is a spiking point neuron, and one popular model is the adaptive exponential integrate-and-fire model (Brette & Gerstner, 2005).

The authors are involved in the Human Brain Project (HBP). To increase reproducibility and enable others to easily use models, the HBP uses a standard set of tools. Nest (Gewaltig & Diesmann, 2007) is commonly used to simulate neurons. PyNN (Davison et al., 2008) is used as middleware to specify the topology and eases the switch from one neural simulator or emulator to another.

Integrate and fire neurons are simple models of neurons, but they are widely used, and can accurately model firing behaviour of biological neurons. They are also computationally efficient to simulate.

One key question about cognitive behaviour is time. It has been noted that many connectionist schemes do not have time in them naturally (Elman, 1990). This is not the case with

simulations of biological neurons as they have a biological time course. Moreover, the time course of the neural behaviour is the same as the time course of the psychological behaviour it produces. So, while there are many connectionist models of associative memory (e.g. (Willshaw, Buneman, & Longuet-Higgins, 1969)), simulated biological neural models of associative memory are needed.

Stroop Model

There are many empirical findings that fall into the category of the Stroop effect (MacLeod, 1991). The cognitive explanation of the effect include horse-racing models based on different processing strength (Cohen, Dunbar, & McClelland, 1990), different perceptual acquisitions (Melara & Algom, 2003) and different selective attention (Roelofs, 2003). Different computational models have been proposed to simulate the Stroop effect based on different cognitive theories. The first connectionist model of the Stroop effect was built in a multi-layer perceptron and trained using supervised learning via back propagation (Cohen et al., 1990). In the same year, Phaf and colleagues developed a selective attentional model for the Stroop effect (Phaf, Van der Heijden, & Patrick, 1990). A more detailed model was built with sub-networks of sensory detection, motor response, attention control and habitual response (Kaplan, Şengör, Gürvit, & Güzeliş, 2007). In those models, the difference in response time was achieved by setting higher distraction on colour naming. A Hopfield network model for the Stroop effect was built and trained with combined patterns of attention and sensory inputs(Yusoff, Grüning, & Browne, 2011). The network converged to trained patterns based on the part-completion results of the attentional modulation in the testing set.

The authors have simulated the Stroop effect with simulated neurons using in Nest with PyNN.[1] Neurons representing ink colour, word, and outputs are modelled with leaky integrate-and-fire neurons. There are CA groups representing ink colour and word, and both are divided into two subgroups representing red or blue conditions. Excitatory connections within a CA spread activation leading to further activation. Inhibitory connections across conditions slow this spread. For instance, neurons in the CA for *blue ink* inhibit neurons in *red word* and vice versa. Synapses from the ink CAs excite their associated output neurons yielding the resulting time.

The response times, see table 1, were different across different conditions in colour-naming, which was interpreted as an interference of voluntary control (MacLeod, 2014).

Semantic Net Model

A plastic neural model of the question answering task (Collins & Quillian, 1969) was developed. The neural model was a variant of the adaptive exponential integrate-and-fire model (Brette & Gerstner, 2005). The learning mechanism was a compensatory Hebbian mechanism, see equations 1 and 2. The training regime was quite precise and the synapses were changed from plastic to static during the simulation, so that they no longer changed after training.

[1]The code for both models is available on http://www.cwa.mdx.ac.uk/NEAL/code/questionICCM.tar.gz.

Table 1: Response time in experiments and simulation

Conditions	Experiments (ms)	Simulation(ms)
Ink incongruent	795	660
Ink congruent	605	436

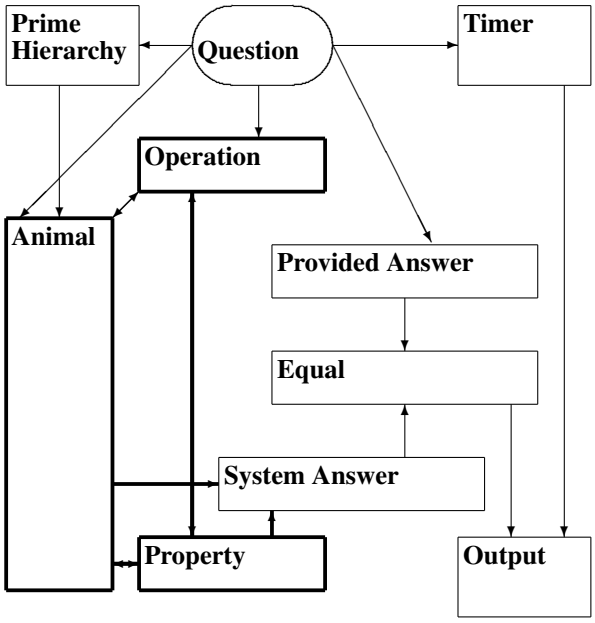

Figure 1: Gross Topology of the Question Answering Associative Memory. Boxes represent sets of neurons. Thick boxes and arrows are plastic. The oval represents the question with spike sources instead of neurons.

A pre-compensatory learning rule was developed in Nest as a synapse model. A pre-compensatory rule targets the total synaptic strength leaving a neuron, W_B in equations 1 and 2. Since the target and current strength need to be stored on the neuron, a new neural model was also developed in Nest. This modified the adaptive exponential integrate-and-fire model by including this constant and variable. The compensatory synapse changed the variable during synaptic weight change.

Figure 1 represents the final neural system. The system consists of three sets of neurons that have learned the semantic net. There is an animal inheritance hierarchy in the animal neurons. Associations are three way between animal, operation and property. For instance, if the *canary is yellow* association is stored, the synaptic weight between animal *canary* and operation *is* has increased weights, as do the synapses between *canary* and property *yellow*, and between *is* and *yellow*.

Training took place in two phases. The first phase learned the animal hierarchy and the second learned the associations.

Initially a well connected net of 200 neurons, the animal neurons, was trained to store 20 concepts and 19 direct hierarchical relations in a three level hierarchy (Animal (Amphibian) (Fish Shark Salmon Bass Pike) (Bird Canary Ostrich Robin Goose Pigeon) (Mammal Dog Cat Rat Bear Monkey Human)). A single concept was represented by 10 neurons with the first five being used for the hierarchy, and the second five used for associations. This was a simple model of a CA.

Neurons were stimulated externally in epochs. Each epoch went through a CA phase, a hierarchy phase and a one neuron firing phase lasting a total of 5000 ms of simulated time. The CA phase stimulated (to firing) each of the 10 neurons in a CA, and then inhibited the entire net after 40 ms. All 20 CAs were stimulated in turn, one after the other 50ms apart. This was followed by the hierarchy phase with each of the 19 direct hierarchical relations presented by stimulating (to firing) the first five neurons of each of the pairs; the entire net was inhibited after 40ms and the next pair was presented. This was followed by a period where each neuron was fired one at a time 11 ms apart. This allowed each synapse to apply the Hebbian decrease rule equation 2. Without this, the synapses between the neurons in the first half of a CA and the synapses between neurons in the second half only go up as they always fire together. This total epoch length was 20x50ms for the CAs, plus 19x50ms for the hierarchical relations, plus 200x11 for the one neuron firing phase, which was rounded to 5000. Hierarchical training took 30 epochs for 150000ms or 150 seconds. In the early epochs neurons fired once during CA and hierarchy presentation, but as the synaptic weight increased, they fired for several times. Inhibition 40ms after presentation allowed the next item to be presented.

The system was run in 1 ms time steps, and the compensatory mechanism considered the neurons co-firing if they fired in the same cycle or with the post-synaptic neuron firing within 10ms of the pre-synaptic neuron. In this manner, a simple CA for each animal concept is learned, and the hierarchical relations are learned.

The animal-animal synapses are saved, synapses with small values are pruned for efficiency, and they are loaded back in as static synapses for the second phase of training. In this phase, there were the 20 animal CAs in 200 neurons, 50 operation neurons, and 50 property neurons. The operation neurons were well connected internally, so that each neuron synapsed with every other neuron, and the property neurons were well connected internally. The second five neurons in each animal CA were well connected with both the operation and property neurons; so each had 100 additional synapses leaving them. Similarly each of the operation and property neurons had connections to those five animal neurons per CA, so each had 100 synapses to the animal neurons. The operation and property neurons connected to each other more sparsely with each connecting to 10 neurons (evenly distributed) in the opposite net.

These three nets were trained in two phases; the first learned the five CAs in the property and operation neurons. The second phase learned the five three way associations. After this the weights were saved.

These saved synaptic weights were loaded into the test system with the low weights pruned for efficiency. During testing, all synapses are static. The question is presented by external simulation. There are two types of question: animal relation property or animal isA animal. Both questions start the (neural) timer. The output neurons are two CAs, one for true and one for false. If the timer completes without the true CA coming on, the timer turns on the false CA, and that is the output. In this case, on means that the CA is firing persistently. The true and false output CAs have mutually inhibitory synapses.

For both types of question, the correct provided answer is stimulated; it fires persistently. Neurons in a CA can fire without causing the circuit to fire persistently; once a CA is firing persistently, it ignited. If the question is an animal relation property one, the animal and operation are stimulated. If there is a property associated, via learning, it becomes active. This then turns on the system answer. The equal net is a set of CAs that are only ignited if both the provided and system answers are firing. If there is an equal answer, it turns on the true output CA. The associations spread directly from the base level animal (canary), but can not spread until the super-level category (bird or animal) is activated. Thus it takes longer to retrieve these associations.

When the animal isA animal question, the second type, is asked, the animal, provided answer and timer are all stimulated. Operation is not stimulated, but the prime hierarchy CA is. This sends extra activation to all of the animals, which supports the spread of activation up the hierarchy. As before, the association is retrieved (or not) by the memory, in this case the animal, turning on the system answer CA.

Following Collins and Quillian (Collins & Quillian, 1969), the property associations are labelled P followed by the level, and the superset relation S followed by the level. The results are shown in table 2. Here S0 refers to the sentence *a pigeon is a pigeon*, and P3 refers to the sentence *a canary has skin*. The false sentences are labelled with False and are *a canary has gills* and *a pigeon is a fish*. The Collins column is the time reported in the paper (Collins & Quillian, 1969).

The associations are retrieved in the correct order by inheritance, but are clearly off time wise. It could easily be argued that the start times are due to input and output processes not accounted for by the model. So, 954ms, for instance, could be added to each of the systems times. Still, the timings are off significantly, with the system's times varying over less than 200ms and the subjects' over almost 500ms.

Discussion

The static Stroop model described above shows that it is reasonably easy to develop a spiking model of a particular associative memory. It is not entirely clear how well this mech-

Table 2: Associative Question Retrieval Time in ms.

Question Type	Collins Answer(ms)	System (ms)
P0	1300	51
P1	1380	61
P2	1460	62
P False	1450	235
S0	1000	46
S1	1170	59
S2	1240	130
S False	1400	235

anism scales. However, the number of neurons should scale linearly to the number of concepts. Similarly, if the associations are stored in synapses, and there are a constant number of associations for each concept, they should be storable in a topology like the sparse topologies of the brain.

While neural models have the advantage of parallelism, the real advantage to using neural systems is that they learn. The question answering model described above makes use of learning. While the Hebbian compensatory learning rule has a degree of biological plausibility, the presentation mechanism and shift from plastic to static neurons is clearly not plausible. One could argue that particular neuro-transmitters turn off plasticity, but the authors feel that is really stretching the metaphor. Instead, we view this model as a step toward more complete ones, and a very early step at that. It is using 10 neurons to represent a concept. The 10 neuron CA would persist indefinitely if not explicitly stopped. Once the neurons have stopped, they do not fire again unless stimulated from the environment, which is clearly biologically unrealistic. The words are stimulated directly from the environment; there is no attempt to read, and there is no attempt to actually ground these words in the environment. While it is clearly incomplete, the system does exhibit some symbolic properties of a semantic net. It also exhibits the right direction of timing for spread of activation of a reasonable cognitive model. It should be relatively simple to improve this so that it more accurately generates these times. This could be done by changing synaptic connectivity, or perhaps moving from a 5-5 CA in the hierarchical structure to a 10-5 CA. The neural parameters could be changed to support slower ignition.

However, improving the model by parameter fitting seems like an unpromising way forward for a significantly better neural associative memory. Instead the model could be improved by simple additions. For instance, the learned portion of the question answering model has no inhibitory neurons. Inhibitory neurons can also be used to reduce overall activation, but also support competition between concepts. The CA for fish is mutually incompatible with the CA for bird, so the two may have mutually inhibitory synapses.

The question answering model has also used synapses for associations. While synapses are clearly involved in associations, larger associated cell assemblies may share neurons. That is, neurons are involved in both CAs. When a CA ignites, neurons in it that are involved in other associated CAs are particularly efficient at priming those CAs and may even lead to their ignition. The authors have explore hierarchical CAs (Huyck, 2007) with shared neurons.

Another problem with the model, and indeed most neural models of learning, is that the neurons that learn are directly stimulated from the environment. Clearly, this is not the case in the brain with at most the sensory neurons being directly stimulated from the environment. Somehow learning must move from the sensory neurons into other areas. Again, the authors have made some progress on this (Huyck & Mitchell, 2014) using a Fatiguing Leaky Integrate and Fire neural model that spontaneously fires when it has not fired recently. Integrating spread into new areas into an associative memory, in addition to increasing biological plausibility, would also address input bandwidth problems of, for example, large neuromorphic machines.

The plastic question answering system only used one learning rule, though it did turn off learning. A better system might take advantage of several learning rules. In particular, the system could benefit from long and short term synaptic modification. The current rule changes the synaptic weight after the next firing, and that weight remains changed (though it of course may be modified again). This is not biologically plausible, as long-term synaptic modification is neither permanent (though it can last for months) nor instantaneous. The authors have explored short term dynamics, and hope to continue in the associative memory context. The long-term firing and synaptic dynamics need to be explored so the associative memory both stores old memories, and learns new memories.

Finally, the presentation mechanism in both models was to merely turn on stored neurons that represent symbols. This really is just a different way of using symbols. If the system could learn concepts from interaction with the environment, there would be scope for appreciably more complex concepts; this is the symbol grounding problem, and some progress on learning concepts can readily incorporate mechanisms for closely associating symbols with those concepts.

Conclusion

This paper has presented two neurocognitive models of associative memory. The first uses static synapses and duplicates the timing behaviour of performance on a Stroop task. The second uses a Hebbian compensatory synaptic modification rule to learn a semantic net. Performance on a question answering task is similar to behaviour of human subjects. Both models are implemented in leaky integrate and fire neurons.

These two models are simple, but it is hoped that they are just two early steps in the development of a more sophisticated neural associative memory mechanism. These models can be extended by the use of inhibitory neurons, supporting competition between CAs; associations including shared neurons, supporting a range of degrees of association; the use of neural models that support spread of CAs beyond neurons

that are directly activated by the environment, allowing the neural system to learn to use neurons that are not directly stimulated by the environment; and the combined use of multiple synaptic modification rules, providing improved flexibility with learning and more biological accuracy.

There will be two main strands in task development: symbolic bootstrapping and symbol grounding. Symbolic bootstrapping can use existing or newly developed symbolic semantic nets. These encodings can be learned by a neural system, and new associations can be learned by, for instance, interpreting text. Large semantic nets can be learned in large neuromorphic systems with millions of neurons, which can support exploration of CA and association dynamics.

Symbol grounding will be used for agents (virtual and robotic) that perform tasks. Initial bootstrapped semantic nets may provide memory, but new concepts and associations will be learned from the environment. This will address one of the key problems of AI.

The goal is to generate a substantially better neural associative memory. This memory will be evaluated on, among other things, the Stroop task, and the question answering task.

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).

References

Anderson, J., & Bower, G. (1973). *Human associative memory*. J. Wiley & Sons.

Bi, G., & 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., & Gerstner, W. (2005). Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *J. Neurophysiol.*, *94*, 3637–3642.

Brette, R., Rudolph, M., Carnevale, T., Hines, M., Beeman, D., Bower, J., ... Destexhe, A. (2007). Simulation of networks of spiking neurons: A review of tools and strategies. *Journal of Computational Neuroscience*, *23*, 349–398.

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

Cohen, J., Dunbar, K., & McClelland, J. (1990). On the control of automatic processes: a parallel distributed processing account of the stroop effect. *Psychological review*, *97*(3), 332.

Collins, A., & Quillian, M. (1969). Retrieval time from semantic memory. *Journal of verbal learning and verbal behavior*, *8(2)*, 240–247.

Davison, A., Brüderle, D., Eppler, J., Muller, E., Pecevski, D., Perrinet, L., & Yqer, P. (2008). PyNN: a common interface for neuronal network simulators. *Frontiers in neuroinformatics*, *2*.

Elman, J. (1990). Finding structure in time. *Cogntivie Science*, *14(2)*, 179–211.

Gewaltig, M., & Diesmann, M. (2007). NEST (neural simulation tool). *Scholarpedia*, *2*(4), 1430.

Hebb, D. O. (1949). *The organization of behavior: A neuropsychological theory*. J. Wiley & Sons.

Huyck, C. (2007). Creating hierarchical categories using cell assemblies. *Connection Science*, *19:1*, 1–24.

Huyck, C., & Mitchell, I. (2014). Post and pre-compensatory Hebbian learning for categorisation. *Computational Neurodynamics*, *8:4*, 299–311.

Huyck, C., & Passmore, P. (2013). A review of cell assemblies. *Biological Cybernetics*, *107:3*, 263–288.

Kaplan, G. B., Şengör, N., Gürvit, H., & Güzeliş, C. (2007). Modelling the stroop effect: A connectionist approach. *Neurocomputing*, *70*(7-9), 1414–1423.

Lansner, A. (2009). Associative memory models: from the cell-assembly theory to biophysically detailed cortex simulations. *Trends in neurosciences,*, *32(3)*, 178–186.

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

MacLeod, C. (2014). The stroop effect. *Encyclopedia of Color Science and Technology*, 1–6.

Melara, R., & Algom, D. (2003). Driven by information: a tectonic theory of stroop effects. *Psychological review*, *110*(3), 422.

Oja, E. (1982). A simplified neuron model as a principal component analyzer. *Journal of Mathematical Biology*, *15*, 267–273.

Phaf, R., Van der Heijden, A., & Patrick, T. (1990). Slam: A connectionist model for attention in visual selection tasks. *Cognitive psychology*, *22*(3), 273–341.

Quillian, M. (1967). Word concepts: A theory of simulation of some basic semantic capabilities. *Behavioral Science*, *12*, 410-30.

Roelofs, A. (2003). Goal-referenced selection of verbal action: modeling attentional control in the stroop task. *Psychological review*, *110*(1), 88.

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

Willshaw, D., Buneman, O., & Longuet-Higgins, H. (1969). Non-holographic associative memory. *Nature*, *222*, 960–962.

Yusoff, N., Grüning, A., & Browne, A. (2011). Modelling the stroop effect: Dynamics in inhibition of automatic stimuli processing. In *Advances in cognitive neurodynamics (ii)* (pp. 641–645). Springer.

Zenke, F., Agnes, E., & Gerstner, W. (2015). Diverse synaptic plasticity mechanisms orchestrated to form and retrieve memories in spiking neural networks. *Nature communications*, *6*, 6922.

Zenke, F., & Gerstner, W. (2017). Hebbian plasticity requires compensatory processes on multiple timescales. *Phil. Trans. R. Soc. B*, *372(1715)*.