

Middlesex University Research Repository

An open access repository of

Middlesex University research

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

Ji, Yuehu, Gamez, David ORCID: <https://orcid.org/0000-0002-3075-655X> and Huyck, Christian R. (2018) A brain-inspired cognitive system that mimics the dynamics of human thought. In: AI-2018 Thirty-eighth SGAI International Conference on Artificial Intelligence, 11-13 Dec 2018, Cambridge, UK. (doi:10.1007/978-3-030-04191-5_4)

Final accepted version (with author's formatting)

This version is available at: <http://eprints.mdx.ac.uk/24994/>

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>

A Brain-inspired Cognitive System that Mimics the Dynamics of Human Thought

Yuehu Ji¹, David Gamez¹ and Christian Huyck¹

Middlesex University, London NW4 4BT UK

`c.huyck@mdx.ac.uk`

<http://www.cwa.mdx.ac.uk/chris/chrisroot.html>

Abstract. In recent years, some impressive AI systems have been built that can play games and answer questions about large quantities of data. However, we are still a very long way from AI systems that can think and learn in a human-like way. We have a great deal of information about how the brain works and can simulate networks of hundreds of millions of neurons. So it seems likely that we could use our neuroscientific knowledge to build brain-inspired artificial intelligence that acts like humans on similar timescales. This paper describes an AI system that we have built using a brain-inspired network of artificial spiking neurons. On a word recognition and colour naming task our system behaves like human subjects on a similar timescale. In the longer term, this type of AI technology could lead to more flexible general purpose artificial intelligence and to more natural human-computer interaction.

Keywords: Spiking Neural Network · Small-World Topology · Cell Assembly · Stroop Effect

1 Introduction

In recent years there has been an explosion of interest in AI. This has partly been the result of developments in deep neural networks, which have achieved success in a wide range of areas, such as face recognition [40] and games [37, 31]. AI systems are also capable of doing impressive feats of natural language processing - for example, IBM's recent successes with Watson [8] and Project Debater [22]. These systems have generated a lot of excitement as well as doom-laden predictions about widespread job losses and apocalyptic takeovers by malevolent machine intelligence.

The recent successes in AI have occurred in situations where a large amount of data is available or there is a highly constrained environment. These AI systems are poor at finding new solutions to problems and they do not think in a human-like way. This creates problems when humans have to interact with the AI. Humans typically use their own minds to model the minds of other people, but this breaks down when humans try to understand an AI system, whose behaviour might be controlled by the processing of vast quantities of text. The mind of an artificial intelligence is opaque to the human mind, which makes it difficult for

humans to collaborate with AI systems and teach them new things. It also leads to major issues with transparency, accountability and trust.

Neuroscience has made big advances in recent years. We have a great deal of information about how the brain works and can simulate hundreds of millions of neurons in real time [9]. Models have been built of the fruit fly brain [1], worms [41] and there are ongoing attempts to simulate the human brain with increasing accuracy [29]. These networks not only aim to build an 'in silico' duplicate of a static brain, but also to simulate the learning and developmental processes: capturing how groups of neurons evolve to perform highly complex cognitive functions.

The human brain is the best example of a general-purpose intelligence that we have. So one way of addressing AI's current limitations could be to build brain-inspired systems that 'think' in a similar way to the human brain. A system that works in a similar way to humans could be more easily understood by humans, which would help to address the issues of trust, accountability and transparency.

In our research we are investigating how cognitive systems can be built using brain-inspired spiking neural networks. The basic unit for these models is a cell assembly [21]: a group of neurons that displays persistent activation. Groups of cell assemblies can be wired together, potentially using learning, into brain-inspired cognitive systems. One of our aims is to produce systems that think in a human-like way. A good way of evaluating this is to measure the system using tests that have been developed by experimental psychology. For example, in this paper we describe a system that shows a similar Stroop effect to human subjects. The construction of brain-inspired systems can also help us to understand the human brain. This can lead to a positive feedback loop in which the results from neuroscience and experimental psychology help us to build AI systems and these AI systems lead to better explanations in neuroscience and experimental psychology.

This paper describes an AI system that performs a word recognition and colour naming task in a human-like way on a similar time scale to humans. It is based on a brain-inspired architecture and implemented using spiking neurons. To evaluate the extent to which our system thinks in a human-like way we measured the timing of its word recognition and colour naming when the colour of the word was congruent with its meaning (for example, 'red' written in red ink) and when the colour of the word was incongruent with its meaning (for example, the word 'red' written in blue ink). When humans perform this task there is a well known interference effect, known as the Stroop effect [38], such that human reaction times vary between the congruent and incongruent tasks. We were hoping to reproduce this Stroop effect in our system.

The first part of the paper gives some background on the word recognition and colour naming tasks and the Stroop effect that is observed when humans perform them. The background section also covers previous computer simulations of the Stroop effect and some of the earlier cognitive systems that have been built with cell assemblies. Section 3 describes how we constructed our system and Section 4 gives the results of our experiments and compares the behaviour

of our system with that of human subjects. The paper concludes with a discussion and our plans for creating AI systems that use learning to develop complex cognitive functions.

2 Background

2.1 Word Recognition and Colour Naming Stroop Effect

In word recognition and colour naming tasks the subjects are presented with a colour word, such as 'red' or blue', which is written in coloured ink. In the congruent situation, the colour of the ink matches the meaning of the word (for example, 'red' written in red ink). In the incongruent situation the colour of the word is different from the colour of the ink (for example, the word 'red' written in blue ink). The subjects have to recognize and repeat the word (WR) or name the colour of the ink (CN). When humans perform these tasks they have a faster reaction time in WR tasks compared with CN tasks. Subjects also have slower reaction times on CN tasks in the incongruent situation where the word and ink colour disagree, but the difference in reaction time is not significant in incongruent WR tasks (see Table 1). This shows that word-reading interferes with colour-naming but colour-naming does not significantly interfere with word-reading. This difference in response times is known as the Stroop effect [38].

The Stroop effect is one of the most studied phenomena of cognitive interference and many variations on Stroop's original experiments have been carried out. For example, Glaser and Glaser [13] introduced stimulus onset asynchrony (SOA) when presenting colour and word stimuli at different temporal positions; the compensated processing time for colour-naming did not result in interference on reading words. Dunbar and Macleod [6] modified the words in different rotations so that participants took longer to read the word. In this situation the hindered word reading still interfered with colour naming. Macleod, Colin and Dunbar [27] introduced shape naming and associated different shapes to different colours. After intensive practicing, slower shape naming interfered with faster colour-naming tasks when participants were asked to name the colour of a shape.

A number of theories have been proposed to explain the Stroop effect [26]. Stroop [39] assumed that people can read words much faster than name colours and attributed the interference on colour-naming to incomplete inhibition on the faster processed word-reading. This model was challenged when slow word reading was shown to strongly interfere with fast colour naming. A different theory was put forward by Cohen, Dunbar and McClelland [4], who suggested that there are parallel distributed processes of word-reading and colour-naming. The strength of the word-reading pathway is stronger than the colour-naming pathway, as there are more experiences of reading words. The pathway strength affects the processing speed and priority in the interference scenario. Melara and Algom [30] suggested that there could be perceptual biases towards word reading compared with colour naming. The attention selection is biased by the dimensional imbalance and uncertainty in linking the stimulus to word or colour per-

ceptions. As words are more salient than colours, word-naming would interfere with colour naming. Roelofs [35] argues that there are architectural difference in the processing of colours and words. In this proposal, colour-naming requires more steps than word reading, so colour-naming requires more attention and is less automatic than word reading.

2.2 Computer Simulations of the Stroop Effect

One of the most influential computer simulations of the Stroop effect was built by Cohen et al. [4]. A multilayer perceptron (MLP) network was trained on control conditions when only colour or word inputs were present. The input layer had input nodes for ink colour, word text and task and it was connected to the output layers via a hidden layer. The activation of sensory inputs and commands were fed through the network to produce output vectors. Training in word or colour conditions increased the strength of the pathways. As reading is highly practiced, there were ten times as many word training items than colour training items. The trained MLP was run repeatedly to generate the output time, which was determined when the accumulated output of one of the output vectors exceeded a response threshold. The system was able to reproduce the basic Stroop effect and predicted a reverse-Stroop effect that has been recently observed [45].

Based on the architecture of the Cohen model, Laeng et al. [25] replaced the single unit colour inputs with three input nodes that worked in combination to represent a much wider range of colours. Kaplan et al. [23] introduced a number of models of neuroanatomic components: a habitual response module, an attentional module, an inhibition module and an error detection module. The authors claimed that the multi-perception network was a duplicate of prefrontal circuits. Benbassat and Henik [2] applied an evolutionary algorithm to the Cohen model to generate the outputs. Different parameters of the networks were imported as part of the genome segments and the system evolved based on its performance in control conditions when only colour or word input was present. This network was able to produce the basic Stroop effect. Kello [24] implemented a simplified articulation process with an inhibition model (gain) to control the information flow through the system. By varying response latencies, a different stimulus onset asynchrony effect was simulated. Phaf [34] built representations of experimental conditions as different feature attributes. By manipulating connection weights, the multi-layer perceptron system was able to perform winner-take-all filtering on input patterns. The recurrent connection between different modules forced the system to converge to a stable state for any given input. This system was able to demonstrate the cooperative and competitive interferences that are detected in some of the Stroop tests.

Roelofs [35] simulated the Stroop effect in a interactive activation model [36] of word production. Based on similar architecture, Van Maanen et al. [43] focused on semantic interference during retrieval. Although colour-naming and word-reading had similar input structure and the symbolic processing unit was biased towards words during retrieval, the system was able to duplicate results

for semantic gradient effects [14] and the SOA effects. Fennell [7] further parameterized the decision process by introducing both boundary and diffusion models [15]. Yusoff, Gruning and Browne [46] simulated the Stroop tests in an auto-associative Hopfield network. This network was initiated with partially enabled (turned on) bits forming colour or word memory as task requirements, and the system took different numbers of cycles to converge on a stable state under different conditions.

2.3 Cognitive Models with Cell Assemblies

The Cell Assembly hypothesis is that CAs are the neural basis of internal representations of concepts, ideas and mental states [16]. A CA is a group of neurons that has relatively high synaptic connectivity, and relatively highly weighted synaptic connectivity. Thus, once some neurons within a given CA start to fire, there is a cascade of firing that causes the larger portion of the CA population 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. When a concept (for example, a word red) is presented to a participant, a group activation pattern of neurons will emerge and persist when a 'red' word is recognized by the participant.

One of the core principles of Hebbian theory [16] is that CAs emerge from Hebbian plasticity; if a group of neurons often fire together, their connecting weights will increase and the particular co-firing neurons will form a CA that acts as a basic unit for neuronal computation. There is considerable theoretical support for CAs (e.g. [3]) and there is a large community of researchers that, in essence, assumes that the CA hypothesis is correct (see [21] for a review). Although imaging data is consistent with the idea of CAs, it cannot provide conclusive evidence because it is not currently possible to record the spiking behaviour of all neurons [11].

In previous work, CAs have been used to build a number of cognitive systems. Researchers have been developing CA based systems for quite some time [33]. In particular, the authors have done work on, for example, associative memory [18], natural language parsing [19], and category learning [20]. Others have developed, for example, systems for robot control [42], and the semantics of words [10]. 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.

3 Methodology

Our cognitive system was constructed using the spiking neuron simulation platform NEST [12]¹.

¹ The code can be found on <http://www.cwa.mdx.ac.uk/NEAL/NEAL.html>

Integrate and fire neurons were programmatically connected together into eight Cell Assemblies (CAs) (see figure 3). There is no learning in the neural network described in this paper. All of the CAs contained one thousand neurons except for the two word reading CAs, which contained two thousand neurons. A CA group was inter-connected with excitatory synapses that has static and identical weights.

The internal connections in a CA follow a small-world-topology [44]. A small world topology is a sparse topology, where nodes can reach most other nodes in a small number of, in this case, synapses. In this paper, we adopted a small world topology that was consistent with neural biology, in which neurons that have more connections are more likely to be connected, which has been reported with 'hub' neurons that are highly connected to other neurons in brain networks [17]. A demonstration of small world topology is shown in figure 1 that shows a network of 12 neurons with a small world topology (left) and a random topology (right). Notice how some neurons in the small world topology are receiving more synaptic connections while the synaptic connections in the random topology are more evenly distributed.

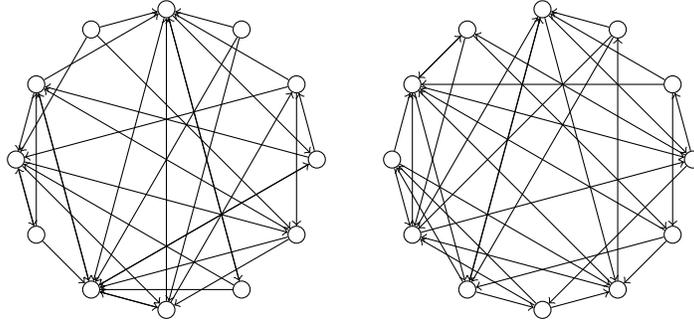


Fig. 1. Small-World Topology.

The small world topology in our project are implemented as follows: initially, each neuron in the CA is connected to its adjacent neuron in a ring. Each connection is then rewired according to equation 1, which shows how likely each of the other neurons is to be selected as the post-synaptic neuron. In this terminology, the connections are one directional synaptic connections. This is a 'rich get richer' policy: post-synaptic neurons that already have many incoming connections have a higher chance of getting a new connection.

$$p_i = \frac{c_i + 1}{\sum_{j=1}^{n \neq m} (c_j + 1)} \quad (1)$$

n is the total number of neurons and $n \times s$ is the total number of uni-directional synaptic connections, with s being the number of synapses leaving each neuron.

p_i is the probability of the i th neuron being selected by a presynaptic neuron whose connection is being rewired. c_i is the number of presynaptic neurons that are currently connecting to the i th neuron. If a connection is to be rewired from neuron m , it will always reconnect to another postsynaptic neuron. Note that we are not implying that small world topologies develop in the brain using this rewiring mechanism, just that they are common topologies in the brain.

One of the properties of a CA is the ability to 'ignite' when some of the neurons in the CA are stimulated. In our experiment, 10 neurons (20 for word CAs) were stimulated repeatedly. The CA will gradually have more neurons starting to fire. A demonstration of CA ignition is shown in figure 2. There are 30 connections leaving each neuron to connect to post-synaptic neurons. It is the same with the random connection condition. Note how the random line never really has any extra firing, while the small world line slowly builds up firing. The random 100 line does increase its firing and rapidly goes to saturation.

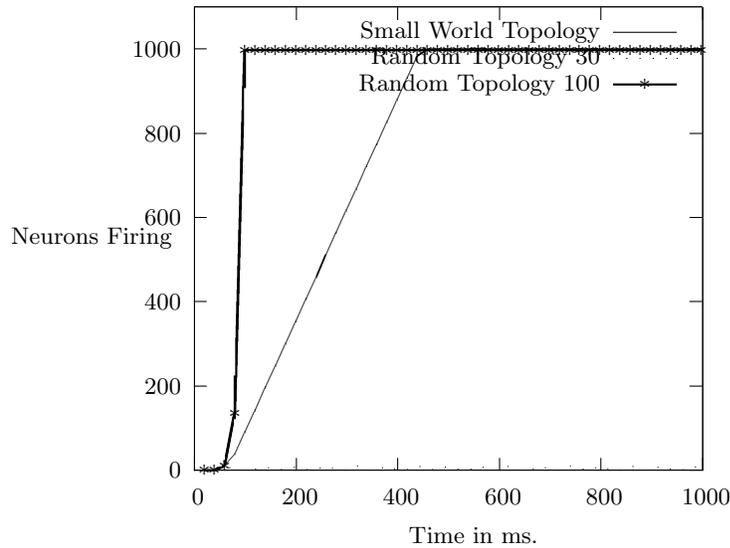


Fig. 2. Firing behaviour of groups of neurons with clamped input, and differing topologies.

The small-world topology was also applied between CAs following the same 'rich get richer' policy under equation 1, however only the count of synapses between these populations was used in the equation. Each neuron in a CA was connected to 1.1% of other neurons in the CA. A gross network structure is shown in figure 3. All CAs had excitatory internal connections, neurons in colour and word CAs had both excitatory and inhibitory external connections. Neurons in task-selection and output CAs only had inhibitory external connections. Excita-

tory connections are shown as solid arrows and inhibitory connections are shown as dotted lines.

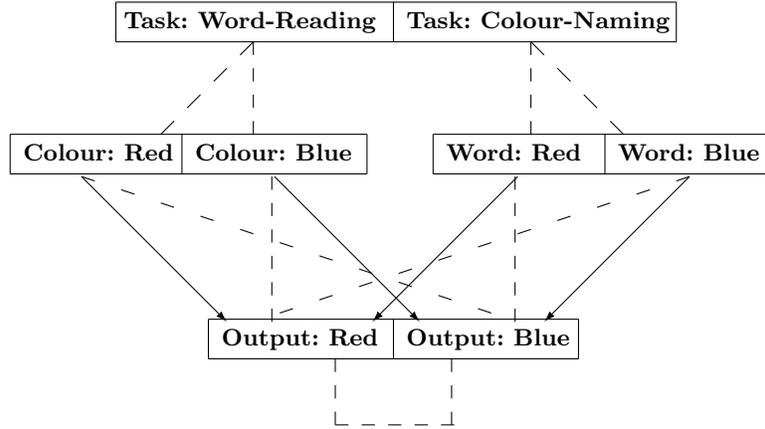


Fig. 3. Network Structure. There are eight Cell Assemblies consisting of neurons with excitatory connections between them. There are also inhibitory connections between CAs, represented by dashed lines, and excitatory connections between CAs, represented by solid lines. All inhibitory connections are from top to bottom except the output CAs are mutually inhibitory.

At the start of a simulation run the task, word and ink CAs were externally activated by injecting a current of 378 mA into a percentage of their neurons. For instance, if it was a word naming (WN) task and the word red was presented in blue ink, then neurons in the word naming task CA, the red word CA and the blue ink CA were activated with the 378 mA injection current, which was continued throughout the simulation of 1000 milliseconds (ms). Activated neurons continued firing throughout the simulation. The percentage of neurons that were externally stimulated for each CA is 0.13% for the ink colour and word CAs, and 0.26% for the task CAs.

The ink colour and word CAs had excitatory connections to the output CAs. Each of these were connected to 0.5% of the appropriate output neurons. They also had inhibitory connections to 0.6% of the appropriate output neurons. The task neurons had connections to 1.2% of the appropriate neurons. (Appropriate here refers to the arcs in figure 3).

The synaptic weights were static and did not change with training or repeated exposure to the task. The synaptic weights were initially set at the minimum weight to evoke spikes in post-synaptic neurons. Then the connecting weights and nodes inside and across different CAs were manually tuned to produce the

firing delays in output neurons that matched the empirical data from human subjects.

Six conditions were simulated, which are shown in Table 1. In the two control conditions either a single word or ink colour CA is stimulated. In the other four conditions, one task, one word and one ink colour CA was stimulated. These four conditions were the congruent or incongruent, colour naming or word reading conditions. So, when one incongruent word reading task was run, the word reading CA, the red ink CA, and the blue word CA were turned on. In this case the correct answer was blue.

4 Results

The empirical data from human subjects and the experimental measurements of our network’s performance are shown in Table 1. The empirical data is a group average of several human participants obtained from Dunbar and Macleod’s work [6]. There are three conditions for each task, for congruent and incongruent conditions refer to a colour-word match and mismatch condition. The control condition is when a stimulus of colour or word was presented without the interference (for colour task was to tell the colour of a coloured box and for word task the participants were presented with words that were not related to colour descriptions).

The reported times for our network are the ramp-up time for the output CA. The ramp-up time of a CA is the moment at which more than 75% of neurons in the CA have fired at least once in the last 10 ms. The system output was recorded when a target CA had ramped-up. If both CAs were ramping-up, the CA that first reached 75% was reported. In isolation, the ramp-up time for the task CAs is 90ms, for colour-naming it is 188ms, and for word-reading it is 146ms, that is faster than colour-naming

Human participants take roughly 300 ms to process visual input and produce a behaviour response. As our system only performs the cognitive task, the time for perception and movement can be discounted. Consequently, Table 1 shows the reported results in brackets, with the results minus 300ms first. The incongruent times match exactly and the others are within 13 ms.

Table 1. Empirical and simulation results

Human Subjects	Congruent (ms)	Incongruent (ms)	Control (ms)
Colour Naming	310(610)	490(790)	330(630)
Word Reading	210(510)	220(520)	215(515)
Simulation	Congruent (ms)	Incongruent (ms)	Control (ms)
Colour Naming	297	490	370
Word Reading	203	220	216

5 Discussion and Future Work

The results in the previous section show that our system successfully executes the word and colour naming tasks with similar timings to human subjects. While the tasks that our system performs are not particularly intelligent, we believe that this type of system could eventually lead to the development of AIs that think in a human-like way and can be more easily understood and trusted by human subjects.

There is no one size fits all recipe for producing an AI system. When large amounts of data are available, deep neural networks are often the best solution. However, deep learning has limitations [28] and other situations have different requirements, such as one shot learning, transparency, symbol manipulation, etc. This paper has demonstrated how a symbol manipulating AI system can be implemented in spiking neurons in a way that is constrained by the brain architecture and measurements of human performance in experimental psychology. In some situations this approach to AI could be a better choice. In the longer term this type of brain-inspired system could prove to be better at flexible general purpose AI than the current state of the art systems.

Although our neural network is a very approximate model of the brain structures that are responsible for naming colours and words, it does have features that might point towards novel explanations of the Stroop effect in the human brain. For example, the small world topology provides important functionality in our CA based system (see Section 3) and our network’s Stroop effect is partly driven by the fact that the number of neurons in the word CAs is twice that of those in the ink CAs.

Humans learn new concepts by modifying synaptic connectivity in the brain. Synaptic connectivity modifications include as long-term potentiation and long-term depression of synapses caused by the spiking patterns of pre and post synaptic neurons [32]. Some of the co-activation patterns are found to follow Hebbian plasticity: if a neuron repeatedly causes another neuron to fire, their connecting weight will tend to be strengthened [16]. In the future, we are planning to add learning rules to our network so that it could develop its topology dynamically by interacting with its environment. Part of this research will be to develop a collection of learning rules that enables the training of simulated neurons to yield desired functional behaviours.

The hope is that a set of learning rules [47] will lead to a stable usable memory system. One cognitive task is a question answering task around semantic nets [5]. This work will be based around encoding known symbolic associative memories into neurons, a form of symbolic boot-strapping. Large scale neural associative memories, run in large neuromorphic systems, can be used in, for example, text mining. Stroop effects should emerge naturally from the behaviour of large associative memories learned in this fashion as a benchmark test of the system.

Beyond this, symbol grounding will be explored, with long term memory units being learned from an environment by a behaving agent; while learning these units, associations will also be learned. These systems will understand

their environment in ways similar to humans, and that understanding will grow as they continue to operate in that environment.

6 Conclusion

This paper has described a system that performs word recognition and colour naming tasks. It was implemented using spiking neurons and it operated in a human-like way on a human-like time scale - displaying a similar Stroop effect to human subjects. In the longer term this type of system could potentially address some of the limitations of the current generation of AI systems. It can also serve as a cognitive model that can help us to understand the brain.

The next step on the authors' plan is to build spiking neural network that can produce CAs automatically instead of hand-wiring. Another potential direction is to build associative memories that would eventually scaffold the capability to conduct word and colour recognitions and then we could re-test the system on the Stroop task.

7 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

1. Arena, P., Patané, L., Termini, P.S.: An insect brain computational model inspired by drosophila melanogaster: simulation results. In: Neural Networks (IJCNN), The 2010 International Joint Conference on. pp. 1–8. IEEE (2010)
2. Benbassat, A., Henik, A.: Examining the stroop effect using a developmental spatial neuroevolution system. In: Proceedings of the Companion Publication of the 2015 Annual Conference on Genetic and Evolutionary Computation. pp. 747–748. ACM (2015)
3. Buzsáki, G.: Neural syntax: cell assemblies, synapsembles, and readers. *Neuron* **68**(3), 362–385 (2010)
4. Cohen, J.D., Dunbar, K., McClelland, J.L.: On the control of automatic processes: a parallel distributed processing account of the stroop effect. *Psychological review* **97**(3), 332 (1990)
5. Collins, A., Quillian, M.: Retrieval time from semantic memory. *Journal of verbal learning and verbal behavior* **8**(2), 240–247 (1969)
6. Dunbar, K., MacLeod, C.M.: A horse race of a different color: Stroop interference patterns with transformed words. *Journal of Experimental Psychology: Human Perception and Performance* **10**(5), 622 (1984)
7. Fennell, A.: Does Response Modality Influence Conflict? Modelling Vocal and Manual Response Stroop Interference. Ph.D. thesis, The Ohio State University (2017)
8. Ferrucci, D., Brown, E., Chu-Carroll, J., Fan, J., Gondek, D., Kalyanpur, A.A., Lally, A., Murdock, J.W., Nyberg, E., Prager, J., Schlaefel, N., Welty, C.: Building watson: An overview of the deepqa project. *AI Magazine* **31**(3), 59–79 (2010)

9. Furber, S.B., Galluppi, F., Temple, S., Plana, L.A.: The spinnaker project. *Proceedings of the IEEE* **102**(5), 652–665 (May 2014). <https://doi.org/10.1109/JPROC.2014.2304638>
10. Garagnani, M., Wennekers, T., Pulvermüller, F.: Recruitment and consolidation of cell assemblies for words by way of hebbian learning and competition in a multi-layer neural network. *Cognitive Computation* **1**(2), 160–176 (2009)
11. Gerstein, G.L., Kirkland, K.L.: Neural assemblies: technical issues, analysis, and modeling. *Neural Networks* **14**(6-7), 589–598 (2001)
12. Gewaltig, M., Diesmann, M.: NEST (NEural Simulation Tool). *Scholarpedia* **2**(4), 1430 (2007)
13. Glaser, M.O., Glaser, W.R.: Time course analysis of the stroop phenomenon. *Journal of Experimental Psychology: Human Perception and Performance* **8**(6), 875 (1982)
14. Glaser, W.R., Dünghoff, F.J.: The time course of picture-word interference. *Journal of Experimental Psychology: Human Perception and Performance* **10**(5), 640 (1984)
15. Gold, J.I., Shadlen, M.N.: The neural basis of decision making. *Annual review of neuroscience* **30** (2007)
16. Hebb, D.O., et al.: *The organization of behavior: A neuropsychological theory*. New York: Wiley (1949)
17. van den Heuvel, M.P., Sporns, O.: Network hubs in the human brain. *Trends in cognitive sciences* **17**(12), 683–696 (2013)
18. Huyck, C.: Creating hierarchical categories using cell assemblies. *Connection Science* **19**:1, 1–24 (2007)
19. Huyck, C.: A psycholinguistic model of natural language parsing implemented in simulated neurons. *Cognitive Neurodynamics* **3**(4), 316–330 (2009)
20. Huyck, C., Mitchell, I.: Post and pre-compensatory Hebbian learning for categorisation. *Computational Neurodynamics* **8**:4, 299–311 (2014)
21. Huyck, C.R., Passmore, P.J.: A review of cell assemblies. *Biological cybernetics* **107**(3), 263–288 (2013)
22. IBM: Project debater, <https://www.research.ibm.com/artificial-intelligence/project-debater/>
23. Kaplan, G.B., Şengör, N., Gürvit, H., Güzeliş, C.: Modelling the stroop effect: A connectionist approach. *Neurocomputing* **70**(7-9), 1414–1423 (2007)
24. Kello, C.T., Plaut, D.C., MacWhinney, B.: The task dependence of staged versus cascaded processing: An empirical and computational study of stroop interference in speech perception. *Journal of Experimental Psychology: General* **129**(3), 340 (2000)
25. Laeng, B., Låg, T., Brennen, T.: Reduced stroop interference for opponent colors may be due to input factors: Evidence from individual differences and a neural network simulation. *Journal of experimental psychology: human perception and performance* **31**(3), 438 (2005)
26. MacLeod, C.M.: The stroop effect. *Encyclopedia of Color Science and Technology* pp. 1–6 (2014)
27. MacLeod, C.M., Dunbar, K.: Training and stroop-like interference: Evidence for a continuum of automaticity. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **14**(1), 126 (1988)
28. Marcus, G.: Deep learning: A critical appraisal. *CoRR* **abs/1801.00631** (2018), <http://arxiv.org/abs/1801.00631>
29. Markram, H.: The human brain project. *Scientific American* **306**(6), 50–55 (2012)

30. Melara, R.D., Algom, D.: Driven by information: a tectonic theory of stroop effects. *Psychological review* **110**(3), 422 (2003)
31. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., and, D.H.: Human-level control through deep reinforcement learning. *Nature* **518**, 529–33 (2015)
32. Nabavi, S., Fox, R., Proulx, C.D., Lin, J.Y., Tsien, R.Y., Malinow, R.: Engineering a memory with ltd and ltp. *Nature* **511**(7509), 348 (2014)
33. Palm, G.: *Neural Assemblies. An Alternative Approach to Artificial Intelligence* (1982)
34. Phaf, R.H., Van der Heijden, A., Hudson, P.T.: Slam: A connectionist model for attention in visual selection tasks. *Cognitive psychology* **22**(3), 273–341 (1990)
35. Roelofs, A.: Goal-referenced selection of verbal action: modeling attentional control in the stroop task. *Psychological review* **110**(1), 88 (2003)
36. Rumelhart, D.E., McClelland, J.L.: An interactive activation model of context effects in letter perception: II. the contextual enhancement effect and some tests and extensions of the model. *Psychological review* **89**(1), 60 (1982)
37. Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., Chen, Y., Lillicrap, T., Hui, F., Sifre, L., van den Driessche, G., Graepel, T., Hassabis, D.: Mastering the game of go without human knowledge. *Nature* **550**, 354–59 (2017)
38. Stroop, J.R.: Studies of interference in serial verbal reactions. *Journal of experimental psychology* **18**(6), 643 (1935)
39. Stroop, J.R.: Factors affecting speed in serial verbal reactions. *Psychological monographs* **50**(5), 38 (1938)
40. Sun, Y., Liang, D., Wang, X., Tang, X.: Deepid3: Face recognition with very deep neural networks. *CoRR* **abs/1502.00873** (2015), <http://arxiv.org/abs/1502.00873>
41. Szigeti, B., Gleeson, P., Vella, M., Khayrulin, S., Palyanov, A., Hokanson, J., Currie, M., Cantarelli, M., Idili, G., Larson, S.: Openworm: an open-science approach to modeling *Caenorhabditis elegans*. *Frontiers in computational neuroscience* **8**, 137 (2014)
42. Tetzlaff, C., Dasgupta, S., Kulvicius, T., Wörgötter, F.: The use of hebbian cell assemblies for nonlinear computation. *Scientific reports* **5**, 12866 (2015)
43. Van Maanen, L., Van Rijn, H.: An accumulator model account of semantic interference in memory retrieval. In: *Proceedings of the Seventh International Conference on Cognitive Modeling*. pp. 322–327 (2006)
44. Watts, D.J., Strogatz, S.H.: Collective dynamics of ‘small-world’ networks. *nature* **393**(6684), 440 (1998)
45. Yamamoto, N., Incera, S., McLennan, C.T.: A reverse stroop task with mouse tracking. *Frontiers in psychology* **7**, 670 (2016)
46. Yusoff, N., Grüning, A., Browne, A.: Modelling the stroop effect: Dynamics in inhibition of automatic stimuli processing. In: *Advances in cognitive neurodynamics (II)*, pp. 641–645. Springer (2011)
47. Zenke, F., Agnes, E.J., Gerstner, W.: Diverse synaptic plasticity mechanisms orchestrated to form and retrieve memories in spiking neural networks. *Nature communications* **6**, 6922 (2015)