Flexibly chaining mental operations with spiking neurons

Hugo Chateau-Laurent[1,2][0000−0002−2891−0503], Chris Eliasmith[2⋆][0000−0003−2933−0209], and Serge Thill[1⋆][0000−0003−1177−4119]

[1] Donders Institute for Brain, Cognition, and Behaviour, Radboud University Nijmegen, Netherlands
hugo.chateaulaurent@student.ru.nl
[2] Centre for Theoretical Neuroscience, Waterloo, ON, Canada

1 Introduction

The ability to flexibly route information between brain regions is crucial to perform new multi-step tasks. The cognitive architecture underlying this mechanism has been investigated with a new paradigm and a series of behavioural experiments [4]. During each trial, participants were shown a digit $N \in \{2, 4, 6, 8\}$ and asked to either compare it to 5 (simple comparison task), add 2 to N then compare the result $I = N \oplus 2$ to 5 (addition-comparison task), or subtract 2 from N then compare the result $I = N \ominus 2$ to 5 (subtraction-comparison task). In each task, participants had to answer "smaller" or "larger". A cycling rule was introduced to operations $\oplus$ and $\ominus$ so that their result stayed within the set of the stimuli. Thus, $8 \oplus 2 \coloneqq 2$ and $2 \ominus 2 \coloneqq 8$. Their results suggest that the human cognitive architecture is suboptimal for chaining operations over concepts like digits due to crosstalk and partially parallel processing. In addition, conscious perception of N was required for chaining, but not for the simple task.

The global neuronal workspace theory of consciousness (see [3] for a recent review) proposes that a network of interconnected neurons breaks the modularity of specialist processors by flexibly broadcasting selected concepts. Here, we propose a spiking neural implementation of the theory that accounts for top-down and bottom-up modulated access to the workspace, ignition and broadcast. We then test the ability of this architecture to perform the tasks.

2 Methods and model

We designed and implemented the model using the Neural Engineering Framework (NEF) [2] and the Semantic Pointer Architecture (SPA) [1]. Following the first principle of the NEF, the activity of a population of neurons represents a vector of real numbers. Using the second principle, a transformation of this vector can be implemented by a connection to another neural population. Connection weights are determined through offline optimisation based on the desired transformation. The NEF contrasts, but is not inconsistent with, approaches where

⋆ These authors contributed equally to this work.

weights are learned. The SPA is a vector-symbolic architecture in the sense that it encodes concepts using a vector representation and combines them with mathematical operations like vector addition for concept superposition and circular convolution for concept binding. Used in combination with the NEF, it provides a framework for modelling high-level cognitive functions in networks of spiking neurons. Cortico-basal ganglia-thalamocortical loops play a key role in this framework and our model.

Figure 1 describes the model architecture. In order to perform the task described in [4], we use 5 processors that communicate exclusively through the global workspace (Figure 1.A ; but see unimplemented crosstalk connections). Each processor is composed of a *Specialist system*, which is a winner-takes-all associative memory (WTA) and a *Receive* population that represents a scalar value (Figure 1.B). The WTA of *Add*, *Subtract* and *Compare* associate their winning input concept to the result of their mathematical operation (e.g. *Add* outputs "FOUR" when its input is "TWO"). *Content* and *Source* are WTA that store the concept being represented in the workspace and the identity of the sending processor, respectively. Unlike *Specialist system*, they amplify their winner to an activation of 1 (ignition). Cortico-basal ganglia-thalamocortical loops gate access to the workspace, such that only the processor with the highest utility sends its content. Access utility of a processor is determined by both the amount of information it stores and its *Attention* level. Similarly, broadcast is gated according to a *Receive* level. *Attention* and *Receive* levels are set by *Router* with rules handcrafted specifically to perform the tasks.

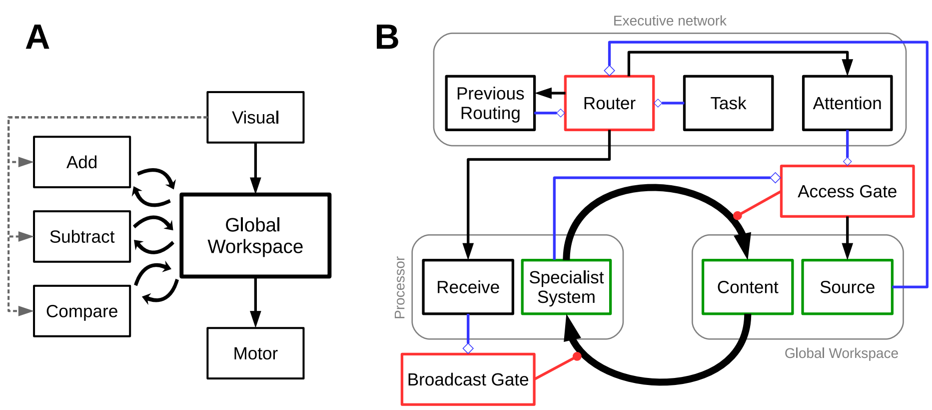

Fig. 1. (A) Overview of the processors and the global workspace. Dashed grey lines indicate possible crosstalk connections described in [4], but are not yet implemented. (B) Interactions between one processor, the global workspace, and the access, broadcast and executive systems. Action selection modules are shown in red, while associative memory networks are in green. Blue connections are used to compute action utilities. While an action is not selected, black excitatory projections are inhibited by red thalamic connections.

3 Results and discussion

The behaviour of the model is shown in Figure 2. A brief visual stimulus "TWO" is presented for 29 ms, following the presentation of a fixation cross. Concepts are sequentially routed from *Visual* to the appropriate operation processors, then to *Motor*. This demonstrates that a global workspace architecture is capable of breaking modularity and execute sequential tasks in spiking neurons working in parallel. However, this is an optimal implementation that ignores crosstalk connections between processors. It is expected that adding such connections will bring more psychological plausibility and allow for a more detailed comparison with the results described in [4].

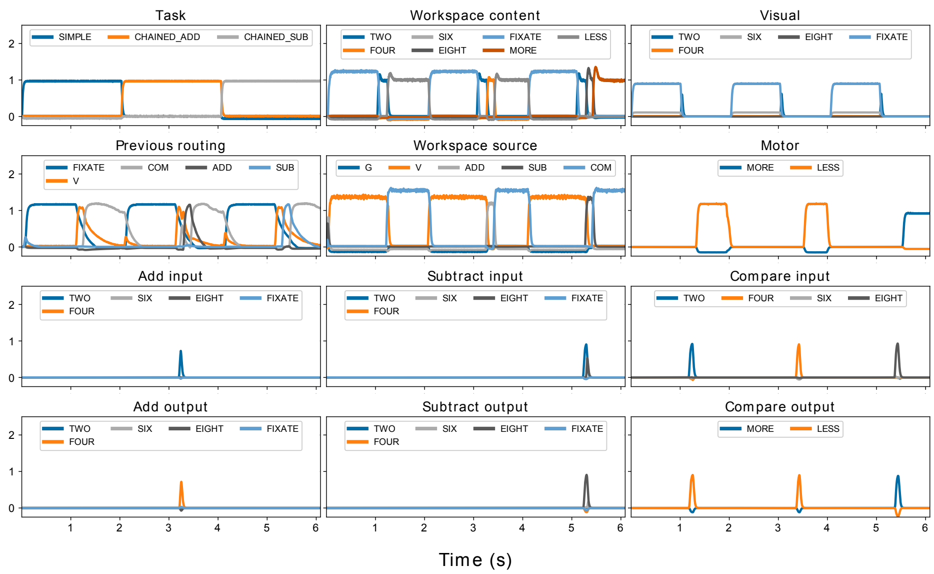

Fig. 2. Behaviour of the model for each task. The plots show similarity of neural activity with concepts over time.

References

1. Eliasmith, C.: How to build a brain: A neural architecture for biological cognition. Oxford University Press (2013)
2. Eliasmith, C., Anderson, C.H.: Neural engineering: Computation, representation, and dynamics in neurobiological systems. MIT press (2003)
3. Mashour, G.A., Roelfsema, P., Changeux, J.P., Dehaene, S.: Conscious processing and the global neuronal workspace hypothesis. Neuron **105**(5), 776–798 (2020)
4. Sackur, J., Dehaene, S.: The cognitive architecture for chaining of two mental operations. Cognition **111**(2), 187–211 (2009)