Spatial computing: utilizing cortical space to dynamically control the cognitive status of neural representations

Abhirup Bandyopadhyay (abhirup.bandyopadhyay@ki.se)
Department of Clinical Neuroscience, K8 Psychology Group, Karolinska Institute
Stockholm, 17177 Sweden

Pawel Herman (paherman@kth.se)
Division of Computational Science and Technology,
School of Electrical Engineering and Computer Science & Digital Futures, KTH Royal Institute of Technology
Stockholm, 100 44, Sweden

Mikael Lundqvist (mikael.lundqvist@ki.se)
Department of Clinical Neuroscience, K8 Psychology Group, Karolinska Institute
Stockholm, 17177 Sweden

Abstract

Working memory (WM) is the ability to transiently store and selectively control a limited set of information in support of higher-order cognition. However, the neural basis of this flexible control is disputed. We recently proposed that cortical network space might be utilized to dynamically update the cognitive status of memory representations and provided experimental evidence for this principle, referred to as spatial computing. Here we implement spatial computing in neural mass models where space is explicitly represented to test its computational feasibility, structural requirements, and generalization capabilities. We demonstrate that distant-dependent like-to-like connectivity and local winner-takes-all-dynamics, both observed in the cortex, are sufficient requirements. In our implementation, spatio-temporal patterns of externally-imposed inhibition dictate where and when information is stored. Thus distinct memory items are encoded in distinct spatial locations, enabling the network to implement selective task-dependent control of WM-representations. The cortical locations of an item can be dynamically updated as their cognitive status change. The imposed, task-dependent inhibition yields a low-dimensional activity pattern independent of item-specific WM information, thus cognitive control generalizes to new patterns. Further, spatial computing can be implemented in WM-networks relying on either persistent activity or synaptic mechanisms. Synaptic mechanisms facilitate intermittent activation of stored items. By imposing a travelling wave of top-down disinhibition distinct items, stored in distinct parts of the network, are activated at different phases of the travelling wave. Together, we demonstrate that low-dimensional dynamics based on utilizing the spatial dimension of cortical space as an information encoding dimension facilitates flexible WM-control and generalization.

Keywords: Spatial computing; Working memory; Cognitive control; Generalization; Neural mass model; winner-takes-all dynamics.

Introduction

Working memory (WM) is a short-term memory sketchpad where memory items are selectively encoded, read out, and deleted when the information may not be relevant (Oberauer, 2002; D'Esposito, Postle, & Rypma, 2000). However, it is still unclear how the brain executes these control mechanisms selectively for different WM items. Recent studies on the temporal neural dynamics of WM (Lundqvist et al., 2016; Lundqvist, Herman, Warden, Brincat, & Miller, 2018) revealed that interactions between bursts of gamma and beta rhythms are correlated with WM control. These interactions have been interpreted as reflecting top-down control, where the gamma bursts associated with elevated spiking encode and maintain WM content, and beta bursts act as the top-down control by inhibiting gamma activities and controlling the access to WM contents (Miller, Lundqvist, & Bastos, 2018; Lewis-Peacock, Drysdale, & Postle, 2015).

Other experimental studies suggest that WM item representations are quite distributed over the cortex and found to traverse across the cortical network (Zokaei, Ning, Manohar, Feredoes, & Husain, 2014). We recently proposed the concept of spatial computing, where gamma-beta interactions utilize network space to selectively control WM activities of distinct items (Lundqvist et al., 2023). In this concept, that rests on distributed representations, the cortical space is seen as an additional information-encoding dimension (Hazy, Frank, & O'reilly, 2007). In other words, the spatial location of certain WM information in the cortical network space encodes the cognitive status of that item such as its temporal order, prioritization, or status as the encoded item or test probe, etc. Controlling the cognitive status of an item corresponds to controlling its cortical location, which does not require knowledge of the actual content or the precise network connectivity responsible for representing the specific WM item. This control is reflected in a low-dimensional embedding on the high-dimensional neural activities (where neurons selective for different WM content have independent activity profiles), and predicts that nearby neurons have shared low-dimensional activity.

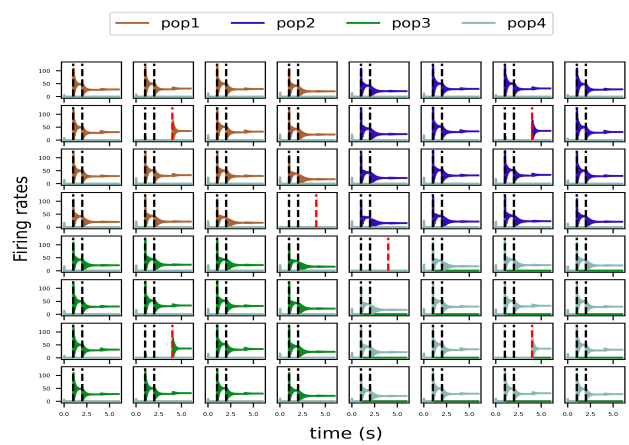

Figure 1: Distance-dependent connectivity allows multiple cell assemblies to be co-active at distinct parts of the network. Patches with the red vertical line are initially inhibited, but following the red line the inhibition is removed. The black and red vertical lines represent the onset of stimulation and disinhibition respectively.

Here we implement spatial computing by imposing low-dimensional spatiotemporal inhibition-disinhibition patterns on the simulated cortical neuron mass model. The spatial component of information encoding is facilitated by distance-dependent connectivity. We demonstrate how this spatial computing-based model allows selective control of task-specific WM items without training the network on the connectivity that helps retain information in WM, and thus generalize to novel items in the spirit of zero-shot learning.

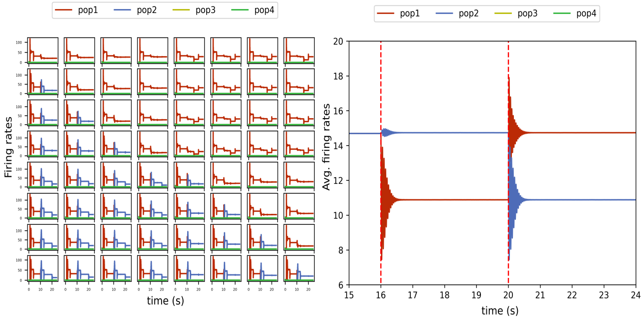

Figure 2: Spatially specific inhibition controls encoding and read-out. A: External stimulation of the whole network of population 1 is followed by stimulation of population 2. Population 2 is only activated in patches where there was brief reset inhibition. B: Between times 16 and 20 top right corner of the network is suppressed, between 20 and 24 the bottom right corner is suppressed.

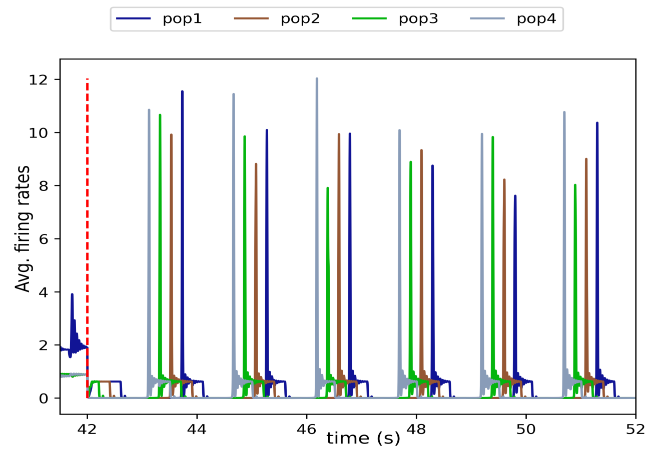

Figure 3: In a network with synaptic augmentation, items can be activated sequentially by spatially specific inhibition. The inhibition is agnostic to the actual populations but works on all populations in a given location.

Results

Here we considered a spatial cortical grid where each cortical patch is represented by a neural mass model that exhibits the winner-takes-all type of dynamics (Feng, Bandyopadhyay, & Mejias, 2023). This spatial grid is connected with distance-dependent excitatory connectivity so that the patches in the close vicinity are strongly coupled. Without distance dependence, a single WM pattern would dominate the whole grid (not shown), but with it multiple patterns than be coactive in the network (Figure 1).

We demonstrate how a change in spatial top-down inhibition can cause a spatial change in how WM content is represented in the cortical space. Strategic inhibition can facilitate storing WM content in selected parts of the network (Figure 1). If some patches are inhibited during the entire stimulation, WM patterns cannot be generated there, only in the patches that are not inhibited. If a previously inhibited patch gets disinhibited the WM pattern can be propagated there, mediated by distance-dependent connectivity, due to the active WM pattern around that patch. Figure 1 describes how distance-dependent connectivity supports updating the spatial location of WM content as an effect of systematic disinhibition (the dotted vertical red line shows the onset of disinhibition) of previously inhibited patches. WM patterns could not be broadcasted to the patches at the center of the grid as there are different patterns activated around it causing strong competition.

For a balanced strength of input stimulus, a new WM pattern could overtake the existing one only if a brief yet strong top-down inhibition (reset inhibition) is provided so that the reset inhibition destroys any previously active pattern. Without reset inhibition, the existing WM pattern remains active even though the entire network receives a new bottom-up stimulus. Figure 2A demonstrates the role of reset inhibition in generating new WM content over the existing one. The entire network is stimulated with first the red population, followed by the blue population. Only the patches below the diagonal receive reset inhibition when the second stimulus arrives. After the patterns are stored in distinct parts of the network they can be controlled independently by up or down-regulating stimulus-independent inhibition spatially (Figure 2B).

By adding synaptic augmentation, the network no longer relies on persistent activity to store WM content (Mongillo, Barak, & Tsodyks, 2008; Lundqvist, Herman, & Lansner, 2011). It was sufficient with periodic disinhibition during which encoded representations were able to briefly become active and refresh the synaptic traces that were responsible for retaining WM content. In this scenario, we again stored different items in different parts of the network. When we imposed a travelling wave (that reaches different parts of the network at different times) of disinhibition it re-activated the different populations at different phases of the wave (Figure 3).

Conclusion

Recent evidence is consistent with the prefrontal cortex implementing spatial computing by distributing WM content and exhibiting spatially organized top-down control in network space (Lundqvist et al., 2023). In this computational study, we provide a framework for such WM control. It was mediated by distance-dependent connectivity and spatial, stimulus-independent patterns of top-down inhibition. This resulted in a low-dimensional embedding of task-related information (Lundqvist et al., 2023). In a nutshell, we developed a novel spatial-computing-based mechanism for selective WM control including encoding WM content, readout, and removing it via the spatial flow of WM representations. The study thus provides a potential explanation for how prefrontal networks may execute flexible yet generalized WM controls via low-dimensional activity in the cortical space.

Acknowledgments

This research was supported by the EU Research and Innovation Funding Programme Horizon 2020 under the project 'Hot-Coal WM', Grant agreement ID: 949131.

References

D'Esposito, M., Postle, B. R., & Rypma, B. (2000). Prefrontal cortical contributions to working memory: evidence from event-related fmri studies. *Executive control and the frontal lobe: Current issues*, 3–11.

Feng, M., Bandyopadhyay, A., & Mejias, J. F. (2023). Emergence of distributed working memory in a human brain network model. *bioRxiv*, 2023–01.

Hazy, T. E., Frank, M. J., & O'reilly, R. C. (2007). Towards an executive without a homunculus: computational models of the prefrontal cortex/basal ganglia system. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *362*(1485), 1601–1613.

Lewis-Peacock, J. A., Drysdale, A. T., & Postle, B. R. (2015). Neural evidence for the flexible control of mental representations. *Cerebral Cortex*, *25*(10), 3303–3313.

Lundqvist, M., Brincat, S. L., Rose, J., Warden, M. R., Buschman, T. J., Miller, E. K., & Herman, P. (2023). Working memory control dynamics follow principles of spatial computing. *Nature Communications*, *14*(1), 1429.

Lundqvist, M., Herman, P., & Lansner, A. (2011). Theta and gamma power increases and alpha/beta power decreases with memory load in an attractor network model. *Journal of cognitive neuroscience*, *23*(10), 3008–3020.

Lundqvist, M., Herman, P., Warden, M. R., Brincat, S. L., & Miller, E. K. (2018). Gamma and beta bursts during working memory readout suggest roles in its volitional control. *Nature communications*, *9*(1), 394.

Lundqvist, M., Rose, J., Herman, P., Brincat, S. L., Buschman, T. J., & Miller, E. K. (2016). Gamma and beta bursts underlie working memory. *Neuron*, *90*(1), 152–164.

Miller, E. K., Lundqvist, M., & Bastos, A. M. (2018). Working memory 2.0. *Neuron*, *100*(2), 463–475.

Mongillo, G., Barak, O., & Tsodyks, M. (2008). Synaptic theory of working memory. *Science*, *319*(5869), 1543–1546.

Oberauer, K. (2002). Access to information in working memory: exploring the focus of attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*(3), 411.

Zokaei, N., Ning, S., Manohar, S., Feredoes, E., & Husain, M. (2014). Flexibility of representational states in working memory. *Frontiers in Human Neuroscience*, *8*, 853.