

were affected differently by muscimol inactivation of medial septum, a manipulation known to result in the break-down of grid cell firing (Brandon et al., 2011). While the firing rate to running speed signal increased in strength, the intrinsic theta frequency/running speed relationship significantly decreased.

Thus, the results suggest that the medial septum provides a speed signal to the grid cells whose interruption during medial septum inactivation causes the breakdown in grid cell firing patterns, and that the critical speed information may be carried by theta rhythmicity rather than firing rate. Accordingly, it may be that coding of speed and direction by subtle changes in burst frequency allows a natural integration into a phase code for location (Burgess, 2008). This model predicts “velocity-controlled oscillators” whose burst frequency has a linear dependence on running speed and cosine dependence on direction, with evidence that “theta cells” found throughout the hippocampal system provide this signal (Welday et al., 2011). This “oscillatory interference” model explains the presence of theta phase precession, and the correlation between grid cell firing and theta rhythmicity during medial septal inactivation (Brandon et al., 2011).

Importantly, in this model, grid cell firing only requires coincident firing caused by the phase coding of the velocity controlled oscillators, and so could generalize to bats and humans in which phase coding may occur in the absence of a strong constant frequency “theta” oscillation.

In continuous attractor models, the bump of activity is shifted by the conjunctive (grid, speed and head-direction) cells to perform translational PI. The existence of these conjunctive cells is powerful evidence for the attractor models, but they do not explain the dependence of grid cell firing on theta rhythmicity during medial septal inactivation, the non-linear dependence of firing rate with running speed (Hinman et al., 2016), or the discrepancy between head-direction and movement direction.

In conclusion, Hinman et al. (2016) provide a fascinating and provocative first insight into the detail of what it is that the medial septum does for grid cells, focusing on speed signaling, and raises plenty of more general questions for future research.

**ACKNOWLEDGMENTS**

We thank the Wellcome Trust for funding this work.

**REFERENCES**

Brandon, M.P., Bogaard, A.R., Libby, C.P., Connerney, M.A., Gupta, K., and Hasselmo, M.E. (2011). *Science* 332, 595–599.

Burgess, N. (2008). *Hippocampus* 18, 1157–1174.

Fuhrmann, F., Justus, D., Sosulina, L., Kaneko, H., Beutel, T., Friedrichs, D., Schoch, S., Schwarz, M.K., Fuhrmann, M., and Remy, S. (2015). *Neuron* 86, 1253–1264.

Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., and Moser, E.I. (2005). *Nature* 436, 801–806.

Hafting, T., Fyhn, M., Bonnevie, T., Moser, M.B., and Moser, E.I. (2008). *Nature* 453, 1248–1252.

Hinman, J.R., Brandon, M.P., Climer, J.R., Chapman, G.W., and Hasselmo, M.E. (2016). *Neuron* 91, this issue, 666–679.

Kropff, E., Carmichael, J.E., Moser, M.B., and Moser, E.I. (2015). *Nature* 523, 419–424.

Sargolini, F., Fyhn, M., Hafting, T., McNaughton, B.L., Witter, M.P., Moser, M.B., and Moser, E.I. (2006). *Science* 312, 758–762.

Taube, J.S. (2007). *Annu. Rev. Neurosci.* 30, 181–207.

Welday, A.C., Shlifer, I.G., Bloom, M.L., Zhang, K., and Blair, H.T. (2011). *J. Neurosci.* 31, 16157–16176.

Yoon, K., Buice, M.A., Barry, C., Hayman, R., Burgess, N., and Fiete, I.R. (2013). *Nat. Neurosci.* 16, 1077–1084.

Zhang, K. (1996). *J. Neurosci.* 16, 2112–2126.

## A Nimble Working Memory

Apoorva Bhandari<sup>1,\*</sup> and David Badre<sup>1,2</sup>

<sup>1</sup>Cognitive, Linguistic, & Psychological Sciences, Brown University, Providence, RI 02912, USA

<sup>2</sup>Brown Institute for Brain Science, Brown University, Providence, RI 02906, USA

\*Correspondence: [apoorva\\_bhandari@brown.edu](mailto:apoorva_bhandari@brown.edu)

<http://dx.doi.org/10.1016/j.neuron.2016.07.030>

In this issue of *Neuron*, Sprague et al. (2016) report fMRI evidence that a degraded working memory representation can be restored by a later cue. The findings raise new questions about the neural mechanisms that underlie such dynamic representational shifts.

Working memory (WM) refers to our ability to maintain and manipulate useful information over short temporal delays (on the order of seconds). It enables us to keep in mind our goal of finding some milk while wandering grocery store aisles, or hold a set of directions in mind as we

plan a driving route. Without the ability to bridge temporal gaps between discrete stimuli and actions with internal representations of information, our cognition and behavior would be much more fragmented and tied to the “here and now.” WM, therefore, is a critical building block

for the construction of intelligent, goal-directed behavior.

A central goal in the study of WM is uncovering the neural mechanisms that support it. A relatively simple mechanism emerged from early studies in non-human primates that identified “memory

cells” in the prefrontal cortex, which showed a selective increase in firing rate to a stimulus that persisted through the delay even after stimulus offset (Fuster and Alexander, 1971). Persistent activity in the WM delay period has also been identified by measurements of the BOLD signal with fMRI (reviewed in Sreenivasan et al., 2014). A wide class of models and theories of memory and cognitive control build on this simple mechanism, implementing WM via actively maintained, stationary neuronal activity patterns for each item in WM (O’Reilly et al., 2010). These models can account for a wide range of phenomena, including the well-known capacity limits of WM. For example, Bays (2015) proposed that WM capacity limits emerge from competitive, inhibitory interactions between neurons coding for the different items in WM. Such competition would lead to reduced sustained firing rates for neurons participating in the population code, thus lowering signal-to-noise and permanently degrading the representation of each individual item.

More recently, however, the simple mechanism of stationary, persistent activity patterns underlying WM has been challenged (Barak et al., 2010; Sreenivasan et al., 2014; Stokes et al., 2013). It has been observed that most individual neurons in the prefrontal cortex do not show sustained firing. Moreover, individual neurons show highly variable tuning profiles which support stable, but highly dynamic coding of information at the population level (reviewed in Sreenivasan et al., 2014; Stokes, 2015). In humans, studies employing fMRI and pattern analysis have found that the information maintained in WM can be decoded from BOLD activity in regions that do not show persistent activity (Sreenivasan et al., 2014). Indeed, as Lewis-Peacock et al. (2012) found, BOLD activity patterns may not even code information maintained in WM if it has not been actively attended. These results are in line with accumulating behavioral evidence that supports a more complex view of WM, in which different bits of information may be differentially accessible (reviewed in Souza and Oberauer, 2016). Indeed, in many situations, the task demands on WM evolve rapidly and it would be adaptive to be able to flexibly shape the

contents of WM. It has been suggested that such flexible, dynamic coding may rely on activity-silent mechanisms of information storage involving, for example, rapid synaptic changes (Stokes, 2015).

Sprague et al. (2016), in this issue of *Neuron*, provide compelling new evidence of such flexibility. They focus on a spate of findings that show that the decline in performance associated with increased WM load can be partially recovered by providing subjects with a retrospective cue (“retro cue”) that signals which item will be subsequently probed (Souza and Oberauer, 2016). How does retro cueing improve performance? One possibility is that the performance benefit of retro cueing is achieved by strengthening the WM representation of the cued target item. This would imply a WM representation that can be dynamically augmented with information maintained elsewhere (in a latent state). Such a finding would challenge a view of WM in which information is primarily maintained via stationary, persistent firing. It would also challenge the aforementioned competitive interaction model of WM capacity limits that predicts a permanent degradation of WM representations in high load conditions (Bays, 2015). On the other hand, it is possible that the retro cue has no impact on the WM representation itself and simply facilitates access or protects it from further decay.

Sprague et al. (2016) addressed this question by directly tracking the fidelity of WM representations during a retro cueing paradigm. Subjects were shown a two-item display and asked to remember the spatial position of one (R1) or both (R2) of the items in WM over a 16 s delay. On the trials when both items were to be remembered, subjects were presented either an informative (R2-valid) or neutral (R2-neutral) retro cue half way through the delay period, which signals the identity of the item whose location would be probed. As expected, Sprague et al. (2016) replicated the WM load effect, finding reduced performance on R2 trials versus R1 trials, and the retro cueing benefit, finding improved performance on R2-valid compared to R2-neutral trials. They then asked if this improvement is accompanied by a strengthening of the underlying WM representation. To assess this, they deployed a model-based image

reconstruction technique. They estimated the weights of a spatial encoding model, which specified the contribution of a set of spatial filters to the activity of each voxel in an ROI, based on data from an independent spatial mapping task. By inverting this model and probing it with delay-period BOLD activity patterns, they were able to generate a moment-by-moment spatial position map that reflected the information about spatial position encoded within the ROI’s activity patterns.

The results are consistent with a dynamic view of WM representations. In the crucial R2-valid condition, prior to the presentation of the retro cue, the estimated spatial maps from both visual and parietal ROIs resembled those from the R2-neutral conditions, representing the spatial locations of both objects. However, after the presentation of the retro cue, the maps of the R2-valid condition began to resemble the R1 map and represented only spatial location of the cued target item. Sprague et al. (2016) quantified the fidelity of the target representation in each condition and showed that the presentation of a retro cue in the R2-valid condition led to an enhancement of the fidelity of the target representation, though it did not approach the fidelity observed in the R1 condition. In other words, they found direct evidence in favor of an enhancement (though not a full recovery) of the cued target in the WM representation. As Sprague et al. (2016) point out, the retro cue itself did not carry any information about the spatial position of the target item, only its identity. Therefore, the additional information about spatial position must have been stored by the brain, even though it was not present in the BOLD activity pattern in the ROIs examined. In other words, given an appropriate cue, latent information in the brain was activated to enhance the active WM representation.

The emerging picture from these results is of a dynamic and nimble WM whose contents are shaped according to current task demands. How might such a cognitive resource be efficiently deployed in the service of task goals? Future work must foreground the control mechanisms that determine, on a moment-by-moment basis, what information is actively coded, held in a latent state, re-activated, or

filtered out of WM. Several computational models have hypothesized a central role for a network incorporating the prefrontal cortex and the basal ganglia in such control (O'Reilly et al., 2010). According to this proposal, the basal ganglia implement input and output “gates” by modulating cortico-thalamo-cortical communication between prefrontal cortex and other parts of the brain. Input gates select stimulus information for encoding into WM. On the other hand, output gates select information from WM and make it available for ongoing cognitive processing, akin to placing it in the “focus of attention” (Souza and Oberauer, 2016). Collectively, these gates allow for a balance between stable maintenance and rapid, selective updating of information in WM. Computational models have tied these WM control dynamics directly to complex, goal-directed behavior (Frank and Badre, 2012).

Important mechanistic questions remain open at the interface of this “gating framework” of WM control and the findings of Sprague et al. (2016). For example, in the gating models, basal ganglia primarily modulate representations in prefrontal cortex. On the other hand, the dynamic modulation of WM representations revealed by Sprague et al. (2016) was local-

ized to visual and parietal cortices. How do these two systems interact with each other? One possibility is that the basal ganglia directly modulate visual WM representations via a parallel set of loops with visual cortex. Alternatively, modulation of visual WM may be mediated by the output-gated, top-down signals from the prefrontal cortex. Indeed, in an analogous task, the presentation of a retro cue is associated with transient activation in a region of the prefrontal cortex, which also increases its coupling with the basal ganglia (Chatham et al., 2014). Moreover, the strength of such coupling has been shown to predict individual differences in response variability, as would be expected from the operation of a stochastic output gate. Therefore, it is possible that a signal from the prefrontal cortex is necessary for the retro cue-mediated recovery observed by Sprague et al. (2016). Future work testing hypotheses like this one will help to flesh out a picture of WM as a dynamic system that incorporates both capacity and its management in the execution of goal-directed behavior.

Sprague et al. (2016) have provided persuasive evidence in favor of a dynamic and nimble working memory system. Their results suggest future experiments that may help identify the control

processes that enable such dynamic coding.

#### REFERENCES

- Barak, O., Tsodyks, M., and Romo, R. (2010). *J. Neurosci.* 30, 9424–9430.
- Bays, P.M. (2015). *Trends Cogn. Sci.* 19, 431–438.
- Chatham, C.H., Frank, M.J., and Badre, D. (2014). *Neuron* 81, 930–942.
- Frank, M.J., and Badre, D. (2012). *Cereb. Cortex* 22, 509–526.
- Fuster, J.M., and Alexander, G.E. (1971). *Science* 173, 652–654.
- Lewis-Peacock, J.A., Drysdale, A.T., Oberauer, K., and Postle, B.R. (2012). *J. Cogn. Neurosci.* 24, 61–79.
- O'Reilly, R.C., Herd, S.A., and Pauli, W.M. (2010). *Curr. Opin. Neurobiol.* 20, 257–261.
- Souza, A.S., and Oberauer, K. (2016). *Atten. Percept. Psychophys.*, Published online April 20, 2016. <http://dx.doi.org/10.3758/s13414-016-1108-5>.
- Sprague, T., Ester, F.E., and Serences, J.T. (2016). *Neuron* 91, this issue, 694–707.
- Sreenivasan, K.K., Curtis, C.E., and D'Esposito, M. (2014). *Trends Cogn. Sci.* 18, 82–89.
- Stokes, M.G. (2015). *Trends Cogn. Sci.* 19, 394–405.
- Stokes, M.G., Kusunoki, M., Sigala, N., Nili, H., Gaffan, D., and Duncan, J. (2013). *Neuron* 78, 364–375.