Theoretical distinction between functional states in working memory and their corresponding neural states

Mark G. Stokes, Paul S. Muhle-Karbe, and Nicholas E. Myers

Wellcome Centre for Integrative Neuroimaging and
Department of Experimental Psychology
University of Oxford

Correspondence to: nicholas.myers@psy.ox.ac.uk or mark.stokes@psy.ox.ac.uk or paul.muhle-karbe@psy.ox.ac.uk

Abstract

Working memory (WM) is important for guiding behaviour, but not always immediately. Here we define a WM item that is currently relevant for guiding behaviour as the functionally 'active' item; whereas items maintained in WM, but not immediately relevant to behaviour, are functionally 'latent'. Traditional neurophysiological theories of WM proposed that content is maintained via persistent neural activity (e.g., stable attractors); however, more recent theories have highlighted the potential role for 'activity-silent' mechanisms (e.g., short-term synaptic plasticity). Given these somewhat parallel dichotomies, it is tempting to associate functionally active and latent cognitive states of WM with persistent-activity and activity-silent neural mechanisms, respectively. In this article we caution against a one-to-one correspondence between functional and activity states. We argue that the principal theoretical requirement for active and latent WM is that the corresponding neural states play qualitatively different functional roles. We consider a number of candidate solutions, and conclude that the neurophysiological mechanisms for functionally active and latent WM items are theoretically independent of the distinction between persistent activity vs activity-silent WM.

Introduction

Working memory (WM) is important for guiding behaviour, but not always immediately. Consider the following scenario: After giving a talk, an audience member asks you a two-part question. As you are preparing to answer question 1, you need to avoid distraction from question 2 lest it interfere with your first answer. However, question 2 still needs to be stored in memory so that you can answer it eventually. Ultimately, both items are equally important. Both need to be robustly encoded and maintained, but only the relevant one should directly influence your current behaviour. Here we define the memory item that is currently relevant as the 'active' item, in the sense that it actively guides ongoing behaviour. By contrast, we define items maintained but not acted upon as 'latent' items, meaning that they should not influence current processing. Latent items have the potential to be brought into an active state once the need arises, but until then, are stored in a robust yet dormant format.

We have previously argued that an active item in WM is maintained in a functionally distinct state from latent items (Myers, Stokes, & Nobre, 2017). In this article, we focus on the variety of neural mechanisms that could plausibly underpin the maintenance of active and latent WM items. Until recently, dominant neural models of WM were based on an unbroken chain of persistent activity that carries mnemonic information during the retention interval (Amit & Brunel, 1997; Camperi & Wang, 1998; Wang, 2001; Wimmer, Nykamp, Constantinidis, & Compte, 2014) . More recently, theorists have proposed that WM could also be maintained via 'activity-silent' neural states, such as short-term synaptic plasticity (Bouchacourt & Buschman, 2019; Manohar, Zokaei, Fallon, Vogels, & Husain, 2019; Zucker & Regehr, 2002). Although these models are not necessarily mutually exclusive, the apparent dichotomy between neurally active and activity-silent mechanisms clearly resembles the functional distinction between active and latent cognitive states of WM (see next section). The purpose of this article is to caution against a direct correspondence between the functional distinction of different memory states and the neural distinction of activity vs activity-silent states. We argue that the principal requirement for distinguishing functionally active from latent WM is that the neural state of active WM should engage with ongoing processing, whereas latent WM should be functionally inert, i.e., not interfere with ongoing processes. Here we consider potential neural solutions for this and outline how they could be tested empirically.

Various Neural Mechanisms Can Support Functionally Latent Working Memory

The distinction between persistently active and activity-silent mechanisms of WM storage has generated vigorous debate (Constantinidis et al., 2018; Lundqvist, Herman, & Miller, 2018), leading some to propose that the two neural mechanisms could serve distinct functions in the service of WM (Masse, Yang, Song, Wang, & Freedman, 2019; Trübutschek, Marti, Ueberschär, & Dehaene, 2019). One popular proposal has been that persistent activity could be associated with attended items in WM, while other items are maintained in a more passive, activity-silent state

(Kamiński & Rutishauser, 2019; LaRocque et al., 2015; LaRocque, Lewis-Peacock, Drysdale, Oberauer, & Postle, 2013; Lewis-Peacock, Drysdale, Oberauer, & Postle, 2012; Manohar et al., 2019; Olivers, Peters, Houtkamp, & Roelfsema, 2011; Stokes, 2015). Although this view has intuitive appeal, the logic that functional WM states should align with this particular dichotomy of activity states has not been thoroughly evaluated in the literature.

In order to fulfil our operational definition, an active WM item should be readily available to interact with ongoing processing, whereas the latent WM item must have minimal influence on processing. Critically, both active and latent items need to be maintained robustly. Further, latent items need to be available to be reformatted into an active state that allows them to effect behaviour when needed.

Functionally active and latent states differentially engage with ongoing processing

We can illustrate these functional criteria with a recent EEG study that measured how active and latent WMs influence WM-guided behaviour (Muhle-Karbe, Myers, and Stokes, 2020). Unlike most WM studies, this experiment was not focused on mechanisms that support maintenance during a retention interval, but rather on the neural dynamics during WM-guided decision-making. Each trial required the discrimination of a visual probe stimulus relative to the 'active' item in WM, but not a 'latent' item. Importantly, the functional status was highly flexible: A latent item could be cued as the relevant item on subsequent trials, making it the active item. By the same token, previously active items could be rendered functionally latent. We found that both active and latent WM items could both be decoded from patterns of EEG activity during the presentation of a probe stimulus. However, only the active item directly affected probe stimulus processing and behaviour. The latent item was represented robustly, but was functionally inactive: it did not influence stimulus processing or behaviour. In this study, WM-based decisions required identifying the difference in angle between an oriented bar held in WM and an oriented probe stimulus. In addition to the contents of WM, the EEG signal also reflected this decision signal, i.e. the angle between a WM orientation and the probe orientation. Even though both WM items could be decoded from the EEG signal, we observed a clear decision signal only relative to the active WM item, but not to the latent item. When the item was in the active state, trial-to-trial variability in decoding predicted the efficiency of WM-based decision-making. By contrast, variability in decoding of items in the latent state only reflected the quality of general maintenance (i.e., the efficacy of the same item once it was converted back into an active state at a later point in time).

Although this study was focused on WM-based decision dynamics rather than on maintenance per se, it does highlight the key functional distinction between active and latent WM states. Below we discuss a number of neural mechanisms that could give rise to the maintenance of WM items for such functionally distinct cognitive states. First, we discuss mechanisms that segregate active and latent items in discrete neural patterns, from separation at the large-scale anatomical level, to separation of activity subspaces within the same neural population, to the frequently

4

invoked separation via neurophysiological mechanism: persistent activity vs activiy-silent states. The key property of all three mechanisms is that that active and latent WM states are statistically uncorrelated (i.e., orthogonal, see also Figure 1), ensuring that active WM states can drive behaviour independently of latent WM states. Finally, we discuss non-independent (i.e., non-orthogonal) coding schemes that have been proposed recently, where there is a statistical dependence between active and latent WM states, and how these might have implications for interference from latent WM on behaviour.

Orthogonal coding schemes for active and latent WM

If active and latent WM have different functional properties, then by definition there must be a difference in the respective neural representation. More specifically, if an item is active in one context, but latent in another context, there must be a distinct coding scheme for the same information in each context (active vs latent; see Figure 1). If they were represented in exactly the same neural state, they would have exactly the same functional properties. Further, we suggest that the difference between active and latent WM states should be *qualitative*, not just *quantitative*. Both active and latent memories need to be maintained robustly: the key difference in not the strength of the coding, but its functional properties. If active and latent neural states are maintained by independent, or orthogonal, coding schemes, they should minimize interference.

Anatomical Separation

Perhaps the simplest means by which latent WM items could be maintained independently is storage in distinct brain areas (Figure 1A, left). For example, a recent fMRI study of visual WM found that active items could be decoded from BOLD activity patterns in a distributed network of brain regions, including early visual, parietal, and prefrontal areas. By contrast, latent items could be decoded only from activity patterns in the intraparietal sulcus and the frontal eye fields (Christophel, Iamshchinina, Yan, Allefeld, & Haynes, 2018), suggestive of a division of labour in coding for active and latent WM items between brain areas. In principle, these findings meet our criterion for qualitatively different coding schemes (Figure 1A, middle and right), since behaviour could be selectively driven by brain areas representing only the active item. However, at the current stage, this proposal is difficult to evaluate, as it would be important to know how these putative WM states actually interact with ongoing processing. Do active WMs affect processing in a task-appropriate manner, while latent WMs remain inert, as in the study by Muhle-Karbe et al. (2020)? Without a clear test of functional significance, it remains possible that brain areas uniquely representing active WM might also be associated with other (theoretically separable and possibly epiphenomenal) processes, such as mental imagery (e.g., Stokes, 2011).

It is also important to note that fMRI findings of decodable delay signals do not strictly imply that WM items are encoded in persistent activity, rather than in 'activity silent' states. Both could in principle result in decodable patterns in BOLD

5

response: Persistent activity should drive statistically separable patterns across voxels, but it is also possible that activity-silent states can be detected in separable patterns of spontaneous activity (Sugase-Miyamoto, Liu, Wiener, Optican, & Richmond, 2008). Given the indirect coupling of spiking activity and BOLD (e.g., Logothetis, Pauls, Augath, Trinath, & Oeltermann, 2001), there are probably even more indirect possibilities that complicate the correspondence between decodable BOLD signals and the underlying neurophysiological mechanism. Nevertheless, whatever the underlying activity state, maintenance in anatomically segregated brain areas fulfils our basic orthogonality requirement.

Separation by different subpopulations (or subspaces) within brain areas

A second possibility is that active and latent WM items are stored within the same brain area(s), or even in an overlapping neural population, without causing interference (Figure 1B). The main prerequisite for this coding scheme is that item-specific activity patterns for the active item are uncorrelated with item-specific patterns when the same item is latent (i.e., statistically separable). This may correspond to latent item patterns occupying the null space of the optimal readout weights for the active item. While the theoretical appeal of such a coding scheme has been noted (Druckmann & Chklovskii, 2012; Spaak, Watanabe, Funahashi, & Stokes, 2017), supporting evidence is still relatively scarce. A recent study (Yoo & Hayden, 2020) found that two stimuli that were both needed for a reward-guided decision were maintained across a delay in orthogonal subspaces of orbitofrontal and ventromedial prefrontal neurons. This separation within the same neural populations could allow downstream brain areas to be driven entirely by one stimulus without interference from the other. A similar mechanism has been demonstrated in movement planning, where premotor cortex maintains a planned movement in a latent state that is 'invisible' to the motor cortex until it needs to be executed (Elsayed, Lara, Kaufman, Churchland, & Cunningham, 2016; Kaufman, Churchland, Ryu, & Shenoy, 2014). Notably, as for the anatomical coding scheme outlined above, such orthogonal patterns could be maintained both via persistent activity, or through activity-independent means (e.g., Hopfield, 1982).

Separation by neurophysiological mechanism

Finally, we consider the proposal that functionally active states are supported by elevated neural activity, whereas functionally latent states correspond to activity-silent mnemonic mechanisms (Kamiński & Rutishauser, 2019; LaRocque et al., 2013; 2015; Lewis-Peacock et al., 2012; Manohar et al., 2019; Olivers et al., 2011; Stokes, 2015). In the framework developed here, this division of labour is only helpful if it confers differential functional properties on active and latent WMs. As highlighted above, it does not bear on the basic maintenance demands: both active and latent memories need to be maintained robustly. Nevertheless, a division of labour between different candidate neurophysiological mechanisms (a stable attractor state based on persistent activity, or short-term synaptic plasticity) could satisfy our key requirement for orthogonal representation (Figure 1C). However, it is also often further assumed that the neurophysiological dichotomy between persistent activity

and activity-silent maintenance naturally aligns with the functional dichotomy between active and latent WM. The intuition seems to be that elevated activity uniquely influences WM-guided behaviour (e.g., via changes in state-dependent processing), and therefore is better suited to active WM, whereas activity-silent mechanisms are effectively functionally dormant. However, it is important to point out that activity-silent mechanisms are not inherently functionally dormant. On the contrary, temporary changes in synaptic connectivity can have a direct functional impact on subsequent processing. For instance, encoding an active item via altered responsivity in the relevant network could allow subsequent input (i.e., a memory probe) to evoke activity that will generate an appropriate response (Stokes et al., 2013), without requiring sustained activity (Manohar et al., 2019; Mongillo, Barak, & Tsodyks, 2008). Therefore, both forms of maintenance (persistent activity, e.g., Mante, Sussillo, Shenoy, & Newsome, 2013; Remington, Narain, Hosseini, & Jazayeri, 2018, and activity-silent connectivity patterns, e.g., Bouchacourt & Buschman, 2019; Manohar et al., 2019) can guide decision-making and support active WM states. Similarly, as we have described above, functionally latent representations could also be maintained via persistent activity as long as they are qualitatively distinct from the corresponding active representation.

It is also important to note that other factors likely determine the extent to which a WM item is associated with persistent activity or activity-silent states. For example, it has recently been proposed that elevated activity could be a signature of current processing or transformation of WM items, rather than storage per se (Masse, Rosen, & Freedman, 2020). If such transformations are more likely to occur on active WM items, this could explain common findings such as active WM being decoded from BOLD activity, while latent items are not (Lewis-Peacock et al., 2012).
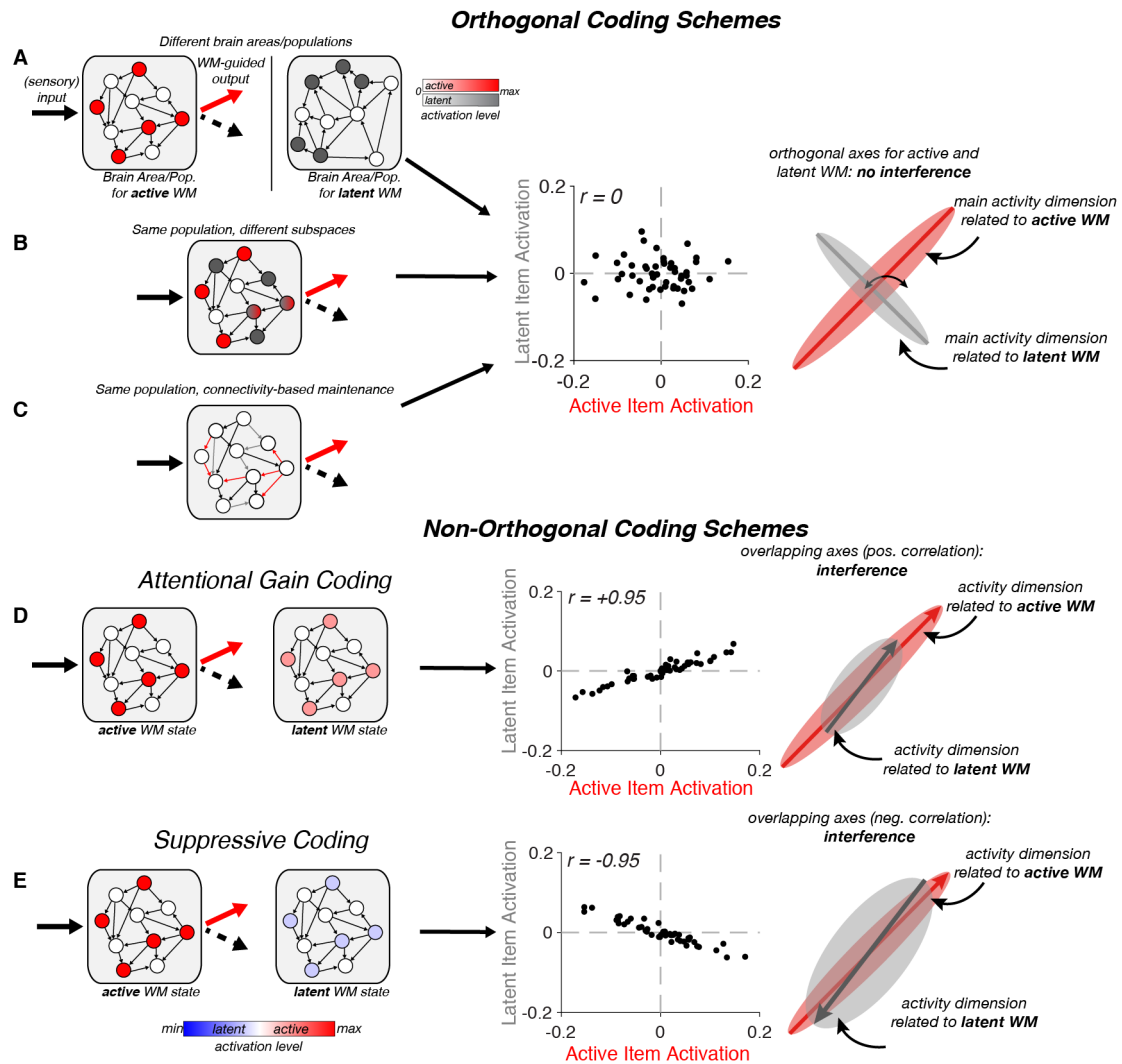
Figure 1. Summary of possible coding schemes for active vs. latent WM. Rows show different putative coding mechanisms for active versus latent WM. Left-hand column: Circuit-level depiction of various coding schemes in an example neural population. Each grey square represents a WM-coding neural population. Within the population, circles represent coding units (neurons), and arrows represent directed connections. Activated units are shown in colour (active: red, latent: grey or blue). Middle column: Correlation between activation patterns for items in an active (x-axis) or a latent (y-axis) state. Individual points indicate units. Correlations are exaggerated for illustration. Right-hand column: Neural state-space representation. When reduced to their most informative dimensions, neural patterns for active or latent items may occupy different subspaces. The extent of their overlap is a reflection of how correlated patterns are for active and latent WM. (A-C). Various coding schemes leading to orthogonal representations (no correlation between active and latent patterns). (A). Separate brain areas or separable neural populations. (B). Separate patterns in the same neural population. (C). Connectivity-based (i.e., activity-silent storage) can also separate active from latent WM by changing the weights of different connections in the population. (D-E). Non-orthogonal coding schemes. (D). Attention Gain coding separates active from latent WM through differences in amplitude, rather than different patterns. (E) Similarly, suppressive coding stores

latent WM in the same neural pattern, but through a reversal of activity, leading to anti-correlated activity patterns that nevertheless occupy the same neural subspace.

Non-orthogonal coding schemes

A number of alternative proposals for neural differences between active and latent items fall into the general category of non-orthogonal coding (Figure 1D-E). For example, Schneegans & Bays (2017) argued that active and latent items are encoded in the same way and differ quantitatively in their level of activation (see also Chun, 2011; Kiyonaga & Egner, 2013). They describe an attractor model that includes an attention parameter that modulates the gain of activity coding for items cued as relevant, compared to latent items (Figure 1D, left). Importantly, persistent activity of items prior to cue presentation is just attenuated, not abolished. If the latent item becomes relevant, activity is increased to the activation level of an active item, allowing more accurate readout. It is worth noting that this model implies that basic maintenance only requires a relatively low energy persistent activity state, whereas the additional activity for attended/active items serves a distinct purpose (e.g., to allow for more reliable readout of the attended item).

In contrast to the orthogonal coding schemes listed above, a difference in activity level between active and latent items means that the underlying patterns are positively correlated (Figure 1D, middle), which presumably could lead to greater cross-talk between active and latent states (Figure 1D, right). For example, the latent item could distort the readout population's estimate of the active item, or create confusion between active and latent items. The severity of this confusion should depend on the relative activation strength of the latent item. One consequence of this should be a trade-off between confusability with the latent item, and more general durability of the memory. While this trade-off could help explain classic WM capacity limits, there is some behavioural evidence that formerly latent items can be restored to an active status with relatively little information loss (Hollingworth & Maxcey-Richard, 2013; Rerko, Souza, & Oberauer, 2014), which seems inconsistent with the model.

Another non-orthogonal coding scheme for active and latent items is suppressive coding. Two recent fMRI studies have suggested that a WM item reverses its activity profile between active and latent states (van Loon, Olmos-Solis, Fahrenfort, & Olivers, 2018; Yu & Postle, 2018). The studies measured BOLD signals in visual and parietal cortex to decode active and latent WM items, and found that when a classifier was trained to discriminate the active WM and was then applied to identify the latent item, its performance dropped below chance level. In other words, item-specific patterns are anti-correlated between the active and latent state. One mechanism to achieve this could be suppression of item-selective neurons when that item enters a latent state (Figure 1E, left and middle).

Such results do not fit neatly within our framework. Even a negative relationship means that a meaningful portion of the activity related to the latent item falls into the subspace coding for the active item (Figure 1E, right). This negative correlation

potentially suffers from the same problem as any mechanism relying on positively correlated patterns between active and latent states: readout trained to discriminate the active item might be influenced by the identity of the latent item. In particular, suppressive coding of the latent item should drive readout of the active item to be *less* similar to the latent one than it actually was, possibly leading to mnemonic repulsion between active and latent items. Interestingly, such repulsion has been reported in certain cases at the behavioural level (e.g., Myers, Chekroud, Stokes, & Nobre, 2018; see also Almeida, Barbosa, & Compte, 2015; Nassar, Helmers, & Frank, 2018). This points to the possibility that suppressive coding could be adaptive in some task environments when similar items need to be disambiguated (e.g., Geng, DiQuattro, & Helm, 2017).

Future Directions
We have laid out a variety of theoretical mechanisms for the storage of functionally active or latent items in WM. Since the key constraint for latent items is that they should not interfere with current behaviour or storage of the active item, the main constraint on possible storage mechanisms is not whether it is neurally active or activity-silent, but rather the orthogonality of the respective coding schemes. Although separation by neurophysiological mechanisms (persistent activity vs activity-silent maintenance) could fulfil this key constraint, it is only one of a larger set of possible solutions. Importantly, it is possible that active and latent items could both be maintained using the same kind of activity state (persistent activity or activity silent maintenance), so long as the mnemonic states are orthogonal: e.g., separate brain areas, overlapping but distinct neural populations, or non-overlapping activity subspaces of the same population.

The distinction between active and latent items echoes the distinction between attentional templates and accessory memory items made by Olivers et al. (2011). The authors distinguished between different neural mechanisms that could underlie the storage of accessory memory items so that it does not drive attentional capture. Our framework can be thought of as an extension of this idea. Attentional capture by stimuli matching the active item (but not the latent item) is one means by which the active item may influence cognition or behaviour. In this instance, the maintenance of the active item would be expected to exert top-down influence on sensory areas so that they preferentially process matching input. In a more general framework, this can be interpreted as one of several possible downstream consequences of the active item that needs to be avoided by the latent item. As we showed recently (Muhle-Karbe et al., 2020), the same principle should also extend to WM-based decision-making (see also Myers, Rohenkohl, Wyart, Woolrich, Nobre, and Stokes, 2015).

Identifying which of the outlined mechanisms support the distinction between active and latent WM requires robust methods for the identification of neural coding mechanisms. However, significant challenges remain for testing candidate neurophysiological coding schemes such as persistent activity and activity-silent coding. In particular, activity-silent states are fundamentally challenging to infer, given that most methods in neuroscience measure some form (or correlate) of

neural activity. Recently, we developed an impulse-response approach to 'ping' activity-silent neural states by measuring the brain's response to task-irrelevant driving input, providing a theoretical potential to track the behaviour of a greater variety of mnemonic states (Wolff, Ding, Myers, & Stokes, 2015; Wolff, Jochim, Akyürek, & Stokes, 2017; see also Rose et al., 2016). While this approach will be useful for enhancing our sensitivity for detecting memory signals that are otherwise undetectable (for whatever reason), it does not strictly adjudicate between the alternative neurophysiological mechanisms. Definitive evidence for activity-silent mechanisms will ultimately require specific evidence of the supposed underlying processes, such as temporary connectivity changes (e.g., periodic refreshing or reactivation of a memory representation, Mongillo et al., 2008), or intrinsic gain modulation (e.g., Stroud, Porter, Hennequin, & Vogels, 2018). At the same time, inferring persistent, uninterrupted activity is not trivial either. Elevated firing during delay periods could reflect transient non-maintenance processes, which can appear to be persistent firing when averaged over many trials (Miller, Lundqvist, & Bastos, 2018).

More generally, it may be insufficient to rely on decodability alone to infer a putative mnemonic state. Decodability has become a ubiquitous marker of WM maintenance (see, e.g., Christophel, Klink, Spitzer, Roelfsema, & Haynes, 2017), but we propose that future work will need to focus less on the mere presence or absence of decodable neural patterns, and more on the functional properties of candidate neural states. This is important because a decodable neural state could be epiphenomenal to WM (e.g., reflect mental imagery or probe expectation, rather than maintenance per se). Moreover, even if the neural state is necessary for WM, it is still critical to understand *how* it influences neural response dynamics to gain a mechanistic understanding of the underlying process (rather than simply identifying the brain area maintaining the WM engram).

Conclusions

In conclusion, we caution against a direct equivalence between functional states in working memory and their corresponding neural states. The key theoretical constraint is that active and latent WMs should be maintained via qualitatively distinct neural states. Within these theoretical constraints, the precise mechanisms of maintenance for either type of WM remains an empirical question. There remain major challenges associated with establishing the neurophysiological mechanisms of maintenance. We argue that focusing on the functional behaviour of putative mnemonic states will be an important future direction.

References

Almeida, R., Barbosa, J., & Compte, A. (2015). Neural circuit basis of visuo-spatial working memory precision: a computational and behavioral study. *Journal of Neurophysiology*, *114*(3), 1806–1818. http://doi.org/10.1152/jn.00362.2015

Amit, D. J., & Brunel, N. (1997). Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cerebral Cortex*, *7*(3), 237–252. http://doi.org/10.1093/cercor/7.3.237

Bouchacourt, F., & Buschman, T. J. (2019). A Flexible Model of Working Memory. *Neuron*, *103*(1), 147–160.e8. http://doi.org/10.1016/j.neuron.2019.04.020

Camperi, M., & Wang, X. J. (1998). A model of visuospatial working memory in prefrontal cortex: recurrent network and cellular bistability. *Journal of Computational Neuroscience*, *5*(4), 383–405. http://doi.org/10.1023/a:1008837311948

Christophel, T. B., Iamshchinina, P., Yan, C., Allefeld, C., & Haynes, J.-D. (2018). Cortical specialization for attended versus unattended working memory. *Nature Neuroscience*, 1–11. http://doi.org/10.1038/s41593-018-0094-4

Christophel, T. B., Klink, P. C., Spitzer, B., Roelfsema, P. R., & Haynes, J.-D. (2017). The Distributed Nature of Working Memory. *Trends in Cognitive Sciences*, *21*(2), 111–124. http://doi.org/10.1016/j.tics.2016.12.007

Chun, M. M. (2011). Visual working memory as visual attention sustained internally over time. *Neuropsychologia*, *49*(6), 1407–1409. http://doi.org/10.1016/j.neuropsychologia.2011.01.029

Constantinidis, C., Funahashi, S., Lee, D., Murray, J. D., Qi, X.-L., Wang, M., & Arnsten, A. F. T. (2018). Persistent Spiking Activity Underlies Working Memory. *Journal of Neuroscience*, *38*(32), 7020–7028. http://doi.org/10.1523/JNEUROSCI.2486-17.2018

Druckmann, S., & Chklovskii, D. B. (2012). Neuronal circuits underlying persistent representations despite time varying activity. *Current Biology : CB*, *22*(22), 2095–2103. http://doi.org/10.1016/j.cub.2012.08.058

Elsayed, G. F., Lara, A. H., Kaufman, M. T., Churchland, M. M., & Cunningham, J. P. (2016). Reorganization between preparatory and movement population responses in motor cortex. *Nature Communications*, *7*(1), 13239–15. http://doi.org/10.1038/ncomms13239

Geng, J. J., DiQuattro, N. E., & Helm, J. (2017). Distractor probability changes the shape of the attentional template. *Journal of Experimental Psychology: Human Perception and Performance*, *43*(12), 1993–2007. http://doi.org/10.1037/xhp0000430

Hollingworth, A., & Maxcey-Richard, A. M. (2013). Selective maintenance in visual working memory does not require sustained visual attention. *Journal of Experimental Psychology: Human Perception and Performance*, *39*(4), 1047–1058. http://doi.org/10.1037/a0030238

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, *79*(8), 2554–2558. http://doi.org/10.1073/pnas.79.8.2554

Kamiński, J., & Rutishauser, U. (2019). Between persistently active and activity-silent frameworks: novel vistas on the cellular basis of working memory. *Annals of the New York Academy of Sciences*, *61*, 267. http://doi.org/10.1111/nyas.14213

Functional vs. neural WM states

Kaufman, M. T., Churchland, M. M., Ryu, S. I., & Shenoy, K. V. (2014). Cortical activity in the null space: permitting preparation without movement. *Nature Neuroscience*, *17*(3), 440–448. http://doi.org/10.1038/nn.3643

Kiyonaga, A., & Egner, T. (2013). Working memory as internal attention: toward an integrative account of internal and external selection processes. *Psychonomic Bulletin & Review*, *20*(2), 228–242. http://doi.org/10.3758/s13423-012-0359-y

LaRocque, J. J., Eichenbaum, A. S., Starrett, M. J., Rose, N. S., Emrich, S. M., & Postle, B. R. (2015). The short- and long-term fates of memory items retained outside the focus of attention. *Memory & Cognition*, *43*(3), 453–468. http://doi.org/10.3758/s13421-014-0486-y

LaRocque, J. J., Lewis-Peacock, J. A., Drysdale, A. T., Oberauer, K., & Postle, B. R. (2013). Decoding attended information in short-term memory: an EEG study. *Journal of Cognitive Neuroscience*, *25*(1), 127–142. http://doi.org/10.1162/jocn_a_00305

Lewis-Peacock, J. A., Drysdale, A. T., Oberauer, K., & Postle, B. R. (2012). Neural Evidence for a Distinction between Short-term Memory and the Focus of Attention. *Journal of Cognitive Neuroscience*, *24*(1), 61–79. http://doi.org/10.1162/jocn_a_00140

Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., & Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature*, *412*(6843), 150–157. http://doi.org/10.1038/35084005

Lundqvist, M., Herman, P., & Miller, E. K. (2018). Working Memory: Delay Activity, Yes! Persistent Activity? Maybe Not. *Journal of Neuroscience*, *38*(32), 7013–7019. http://doi.org/10.1523/JNEUROSCI.2485-17.2018

Manohar, S. G., Zokaei, N., Fallon, S. J., Vogels, T. P., & Husain, M. (2019). Neural mechanisms of attending to items in working memory. *Neuroscience & Biobehavioral Reviews*, *101*, 1–12. http://doi.org/10.1016/j.neubiorev.2019.03.017

Mante, V., Sussillo, D., Shenoy, K. V., & Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*, *503*(7474), 78–84. http://doi.org/10.1038/nature12742

Masse, N. Y., Rosen, M. C., & Freedman, D. J. (2020). Reevaluating the Role of Persistent Neural Activity in Short-Term Memory. *Trends in Cognitive Sciences*, *24*(3), 242–258. http://doi.org/10.1016/j.tics.2019.12.014

Masse, N. Y., Yang, G. R., Song, H. F., Wang, X.-J., & Freedman, D. J. (2019). Circuit mechanisms for the maintenance and manipulation of information in working memory. *Nature Neuroscience*, *22*(7), 1159–1167. http://doi.org/10.1038/s41593-019-0414-3

Miller, E. K., Lundqvist, M., & Bastos, A. M. (2018). Working Memory 2.0. *Neuron*, *100*(2), 463–475. http://doi.org/10.1016/j.neuron.2018.09.023

Mongillo, G., Barak, O., & Tsodyks, M. (2008). Synaptic Theory of Working Memory. *Science*, *319*(5869), 1543–1546. http://doi.org/10.1126/science.1150769

Muhle-Karbe, P. S., Myers, N. E. & Stokes, M. G. (2020). Active and latent states in working memory support decisions at different time scales. *Manuscript in preparation*.

Myers, N. E., Chekroud, S. R., Stokes, M. G., & Nobre, A. C. (2018). Benefits of flexible prioritization in working memory can arise without costs. *Journal of*

Experimental Psychology: Human Perception and Performance, *44*(3), 398–411. http://doi.org/10.1037/xhp0000449

Myers, N. E., Rohenkohl, G., Wyart, V., Woolrich, M. W., Nobre, A. C., & Stokes, M. G. (2015). Testing sensory evidence against mnemonic templates. *eLife*, *4*, 1–25. http://doi.org/10.7554/eLife.09000

Myers, N. E., Stokes, M. G., & Nobre, A. C. (2017). Prioritizing Information during Working Memory: Beyond Sustained Internal Attention. *Trends in Cognitive Sciences*, *21*(6), 449–461. http://doi.org/10.1016/j.tics.2017.03.010

Nassar, M. R., Helmers, J. C., & Frank, M. J. (2018). Chunking as a rational strategy for lossy data compression in visual working memory. *Psychological Review*, *125*(4), 486–511. http://doi.org/10.1037/rev0000101

Olivers, C. N. L., Peters, J., Houtkamp, R., & Roelfsema, P. R. (2011). Different states in visual working memory: when it guides attention and when it does not. *Trends in Cognitive Sciences*, *15*(7), 327–334. http://doi.org/10.1016/j.tics.2011.05.004

Remington, E. D., Narain, D., Hosseini, E. A., & Jazayeri, M. (2018). Flexible Sensorimotor Computations through Rapid Reconfiguration of Cortical Dynamics. *Neuron*, *98*(5), 1005–1019.e5. http://doi.org/10.1016/j.neuron.2018.05.020

Rerko, L., Souza, A. S., & Oberauer, K. (2014). Retro-cue benefits in working memory without sustained focal attention. *Memory & Cognition*, *42*(5), 712–728. http://doi.org/10.3758/s13421-013-0392-8

Rose, N. S., LaRocque, J. J., Riggall, A. C., Gosseries, O., Starrett, M. J., Meyering, E. E., & Postle, B. R. (2016). Reactivation of latent working memories with transcranial magnetic stimulation. *Science*, *354*(6316), 1136–1139. http://doi.org/10.1126/science.aah7011

Schneegans, S., & Bays, P. M. (2017). Restoration of fMRI Decodability Does Not Imply Latent Working Memory States. *Journal of Cognitive Neuroscience*, *29*(12), 1977–1994. http://doi.org/10.1162/jocn_a_01180

Spaak, E., Watanabe, K., Funahashi, S., & Stokes, M. G. (2017). Stable and Dynamic Coding for Working Memory in Primate Prefrontal Cortex. *Journal of Neuroscience*, *37*(27), 6503–6516. http://doi.org/10.1523/JNEUROSCI.3364-16.2017

Stokes, M. G. (2011). Top-down visual activity underlying VSTM and preparatory attention. *Neuropsychologia*, *49*(6), 1425–1427. http://doi.org/10.1016/j.neuropsychologia.2011.02.004

Stokes, M. G. (2015). "Activity-silent" working memory in prefrontal cortex: a dynamic coding framework. *Trends in Cognitive Sciences*, *19*(7), 394–405. http://doi.org/10.1016/j.tics.2015.05.004

Stokes, M. G., Kusunoki, M., Sigala, N., Nili, H., Gaffan, D., & Duncan, J. (2013). Dynamic Coding for Cognitive Control in Prefrontal Cortex. *Neuron*, *78*(2), 364–375. http://doi.org/10.1016/j.neuron.2013.01.039

Stroud, J. P., Porter, M. A., Hennequin, G., & Vogels, T. P. (2018). Motor primitives in space and time via targeted gain modulation in cortical networks. *Nature Neuroscience*, 1–15. http://doi.org/10.1038/s41593-018-0276-0

Sugase-Miyamoto, Y., Liu, Z., Wiener, M. C., Optican, L. M., & Richmond, B. J. (2008). Short-Term Memory Trace in Rapidly Adapting Synapses of Inferior Temporal Cortex. *PLoS Computational Biology*, *4*(5), e1000073.

http://doi.org/10.1371/journal.pcbi.1000073

Trübutschek, D., Marti, S., Ueberschär, H., & Dehaene, S. (2019). Probing the limits of activity-silent non-conscious working memory. *Proceedings of the National Academy of Sciences of the United States of America*, *116*(28), 14358–14367. http://doi.org/10.1073/pnas.1820730116

van Loon, A. M., Olmos-Solis, K., Fahrenfort, J. J., & Olivers, C. N. (2018). Current and future goals are represented in opposite patterns in object-selective cortex. *eLife*, *7*, 1427. http://doi.org/10.7554/eLife.38677

Wang, X.-J. (2001). Synaptic reverberation underlying mnemonic persistent activity. *Trends in Neurosciences*, *24*(8), 455–463. http://doi.org/10.1016/S0166-2236(00)01868-3

Wimmer, K., Nykamp, D. Q., Constantinidis, C., & Compte, A. (2014). Bump attractor dynamics in prefrontal cortex explains behavioral precision in spatial working memory. *Nature Neuroscience*, *17*(3), 431–439. http://doi.org/10.1038/nn.3645

Wolff, M. J., Ding, J., Myers, N. E., & Stokes, M. G. (2015). Revealing hidden states in visual working memory using electroencephalography. *Frontiers in Systems Neuroscience*, *9*(5), 123. http://doi.org/10.3389/fnsys.2015.00123

Wolff, M. J., Jochim, J., Akyürek, E. G., & Stokes, M. G. (2017). Dynamic hidden states underlying working-memory-guided behavior. *Nature Neuroscience*, *20*(6), 864–871. http://doi.org/10.1038/nn.4546

Yoo, S. B. M., & Hayden, B. Y. (2020). The Transition from Evaluation to Selection Involves Neural Subspace Reorganization in Core Reward Regions. *Neuron*, *105*(4), 712–724.e4. http://doi.org/10.1016/j.neuron.2019.11.013

Yu, Q., & Postle, B. R. (2018). Different states of priority recruit different neural codes in visual working memory, 1–35. http://doi.org/10.1101/334920

Zucker, R. S., & Regehr, W. G. (2002). Short-term synaptic plasticity. *Annual Review of Physiology*, *64*(1), 355–405. http://doi.org/10.1146/annurev.physiol.64.092501.114547