REVIEW | *Working Memory: Neural Mechanisms*

Dynamic population coding and its relationship to working memory

Ethan M. Meyers

Center for Brains, Minds and Machines, Massachusetts Institute of Technology, Cambridge, Massachusetts, and School of Cognitive Science, Hampshire College, Amherst, Massachusetts

Meyers EM. Dynamic population coding and its relationship to working memory. *J Neurophysiol* 120: 2260–2268, 2018. First published September 12, 2018; doi:10.1152/jn.00225.2018.—For over 45 years, neuroscientists have conducted experiments aimed at understanding the neural basis of working memory. Early results examining individual neurons highlighted that information is stored in working memory in persistent sustained activity where neurons maintained elevated firing rates over extended periods of time. However, more recent work has emphasized that information is often stored in working memory in dynamic population codes, where different neurons contain information at different periods in time. In this paper, I review findings that show that both sustained activity as well as dynamic codes are present in the prefrontal cortex and other regions during memory delay periods. I also review work showing that dynamic codes are capable of supporting working memory and that such dynamic codes could easily be "readout" by downstream regions. Finally, I discuss why dynamic codes could be useful for enabling animals to solve tasks that involve working memory. Although additional work is still needed to know definitively whether dynamic coding is critical for working memory, the findings reviewed here give insight into how different codes could contribute to working memory, which should be useful for guiding future research.

dynamic coding; neural coding; persistent activity; population decoding; working memory

INTRODUCTION

Working memory refers to the ability to maintain and manipulate a limited amount of information over a period of seconds (Aben et al. 2012; Baddeley 2012; Cowan 2008). Early work examining firing rates of individual neurons found that neurons in the prefrontal cortex (PFC) and other brain regions had "persistent sustained activity," where neurons maintained elevated firing rates during delay periods (Fuster 1973; Fuster and Alexander 1971; Kubota and Niki 1971), and that neurons had "persistent selectivity," where the activity was selective for stimuli that needed to be remembered (Funahashi et al. 1993). Many subsequent studies also found similar persistent sustained activity in the PFC and other brain regions (for reviews, see Leavitt et al. 2017; Riley and Constantinidis 2016). These findings led to the idea that persistent sustained activity is the neural basis of working memory (Nee and D'Esposito 2018), and much work has been devoted to understanding the mechanisms that underlie such sustained activity (Major and Tank 2004; Zylberberg and Strowbridge 2017) and to building computational models that exhibit persistent states

(Brunel 2003; Constantinidis and Wang 2004; Durstewitz et al. 2000; Wang 2001; Wimmer et al. 2014).

Although sustained activity might be the neural basis of working memory, recent studies, often using population analysis methods, have also emphasized that many neurons show patterns of neural activity that are selective for only short periods of time during longer memory delay periods. Thus it is possible that other forms of neural activity, or even synaptic changes (Mongillo et al. 2008; Stokes 2015), are the fundamental mechanisms enabling working memory. In this paper, I review studies that have shown that dynamic neural activity is associated with working memory. I then discuss a few additional results that give further insight into whether dynamic codes could be sufficient for supporting working memory and how such codes could be "readout" by downstream brain regions. Finally, I discuss possible reasons why information in working memory might be stored in dynamic population codes. I focus here on studies that recorded action potentials from individual neurons because these data contain the temporal and spatial resolution to examine neural coding [see D'Esposito and Postle (2015) for a review that focuses more on functional MRI and EEG studies]. I hope the review of these results will help guide future research that will lead to a more precise

Address for reprint requests and other correspondence: E. M. Meyers, Hampshire College, 893 West St., Amherst, MA 02139 (e-mail: emeyers @hampshire.edu).

"algorithmic" explanation for how these types of codes contribute to behavior.

A REVIEW OF DYNAMIC CODING ASSOCIATED WITH WORKING MEMORY

As discussed in the INTRODUCTION, early findings on the neural basis of working memory focused on persistent sustained activity (Fuster 1973; Fuster and Alexander 1971; Kubota and Niki 1971; Miller et al. 1996). Although these studies showed that there was a range of activity profiles in different neurons, such as neurons that responded to the onset of a visual cue (Chafee and Goldman-Rakic 1998; Fuster 1973; Goldman-Rakic 1996), many of these studies highlighted the fact that neurons had tonic firing rates during delay periods (Fig. 1A). One notable exception to the description of tonic delay period activity was discussed in a study by Batuev et al. (1980) that found that most neurons in the lateral PFC of macaques were selective for only shorter periods of time during a memory delay period, leading to what they described as "successive 'relay-race' involvement of different neuronal populations in order to sustain a prolonged activation of a small group of 'summator' units"; however, this paper did not attract much attention.

Around 2000, several studies analyzed single-neuron activity in more detail and began to emphasize neural correlates of working memory that were different from tonic sustained activity (Brody et al. 2003a, 2003b; Romo et al. 1999; Shafi et al. 2007; Warden and Miller 2007; Zaksas and Pasternak 2006). For example, analyses by Zaksas and Pasternak (2006) found that most motion direction neurons in the middle temporal visual area and PFC were only transiently selective during a memory delay period, with most neurons maintaining their selectivity for around 150–300 ms out of the 1,500-ms delay period (Fig. 1B). Likewise, Brody et al. (2003a) found that whereas some PFC neurons were selective to tactile stimuli throughout the delay period, other neurons were only selective for shorter periods of time (Brody et al. 2003a, 2003b; Romo et al. 1999).

Neural population-level analyses also more clearly revealed "dynamic population coding" phenomena associated with working memory tasks that were different from persistent sustained activity (Fig. 2). For example, a study by Baeg et al. (2003) showed that it was possible to decode the direction a rat intended to move during a delay period using only neurons in the PFC that did not have continuous high levels of activity and thus suggested that working memory might be coded by sequential activation of different neurons rather than sustained activity. Additionally, "cross-temporal decoding" analyses, where a pattern classifier is trained with data from one time period and then tested with data from a different time, clearly showed that information was contained in different patterns of activity at different time periods (see Fig. 3 and Barak et al. 2010; Crowe et al. 2010; Meyers et al. 2008, 2012; Nikolić et al. 2009; Pastalkova et al. 2008). For example, Meyers et al. (2008) first used temporal cross-decoding to show there was dynamic coding of category information PFC and the inferior temporal cortex during a working memory task, and Pastalkova et al. (2008) used cross-correlation analyses to show dynamics in rat hippocampus also in a memory delay task. Likewise, Nikolić et al. (2009) and Crowe et al. (2010) used versions of temporal cross-decoding to show dynamic representations in delay periods for activity in V1 and in parietal cortex, respectively. Knowledge about dynamic population coding became more widespread after the publication of Stokes et al. (2013), who used temporal cross-correlation analyses to show dynamic population coding in the PFC during a memory delay in a paired association task, and with the publication of several review papers that discussed dynamic coding and its relationship to working memory (King and Dehaene 2014; Sreenivasan et al. 2014; Stokes 2015).
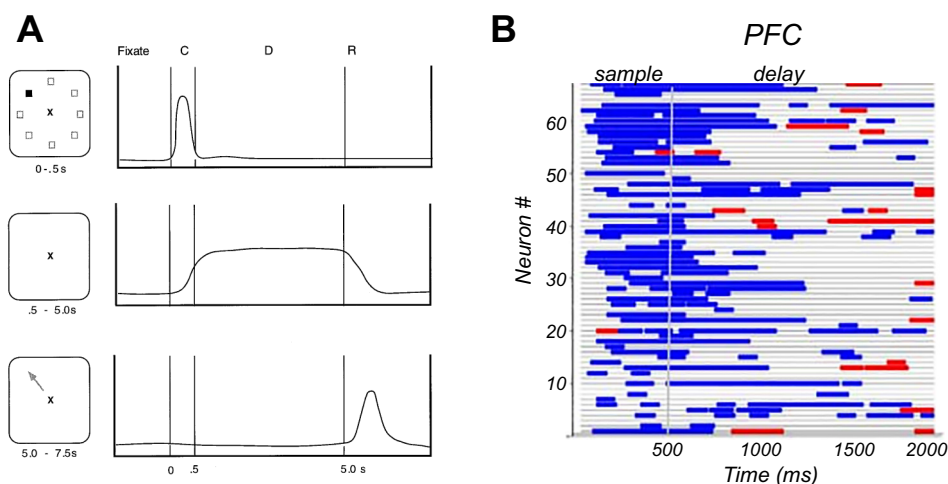

Fig. 1. Single-neuron results emphasizing sustained delay activity and dynamic selectivity. *A*: figure from a review paper by Goldman-Rakic (1996) characterizing different types of neural responses seen in prefrontal cortex (PFC) activity. Results show that neurons respond to cue onset (C; *top* plot) and respond to saccade onset (R; *bottom* plot) and that they often have sustained tonic activity during the delay period (D; *middle* plot). [Reprinted from Goldman-Rakic PS. Regional and cellular fractionation of working memory. *Proc Natl Acad Sci USA* 93: 13473–13480, 1996, with permission. Copyright (1996) National Academy of Sciences, U.S.A.] *B*: analysis of the direction selectivity of 67 neurons in the PFC showing that many neurons are selective for relatively short periods of time during a delay period. Blue regions indicate times at which activity associated with a preferred sample were reliably higher than those associated with an antipreferred sample. Red regions indicate the opposite relationship, and gray regions show time where the signal was not significantly direction selective. Shown is that many neurons only appear selective for short periods of time in the delay period. #, Number. [Results are adapted from Zaksas and Pasternak (2006) with permission.]
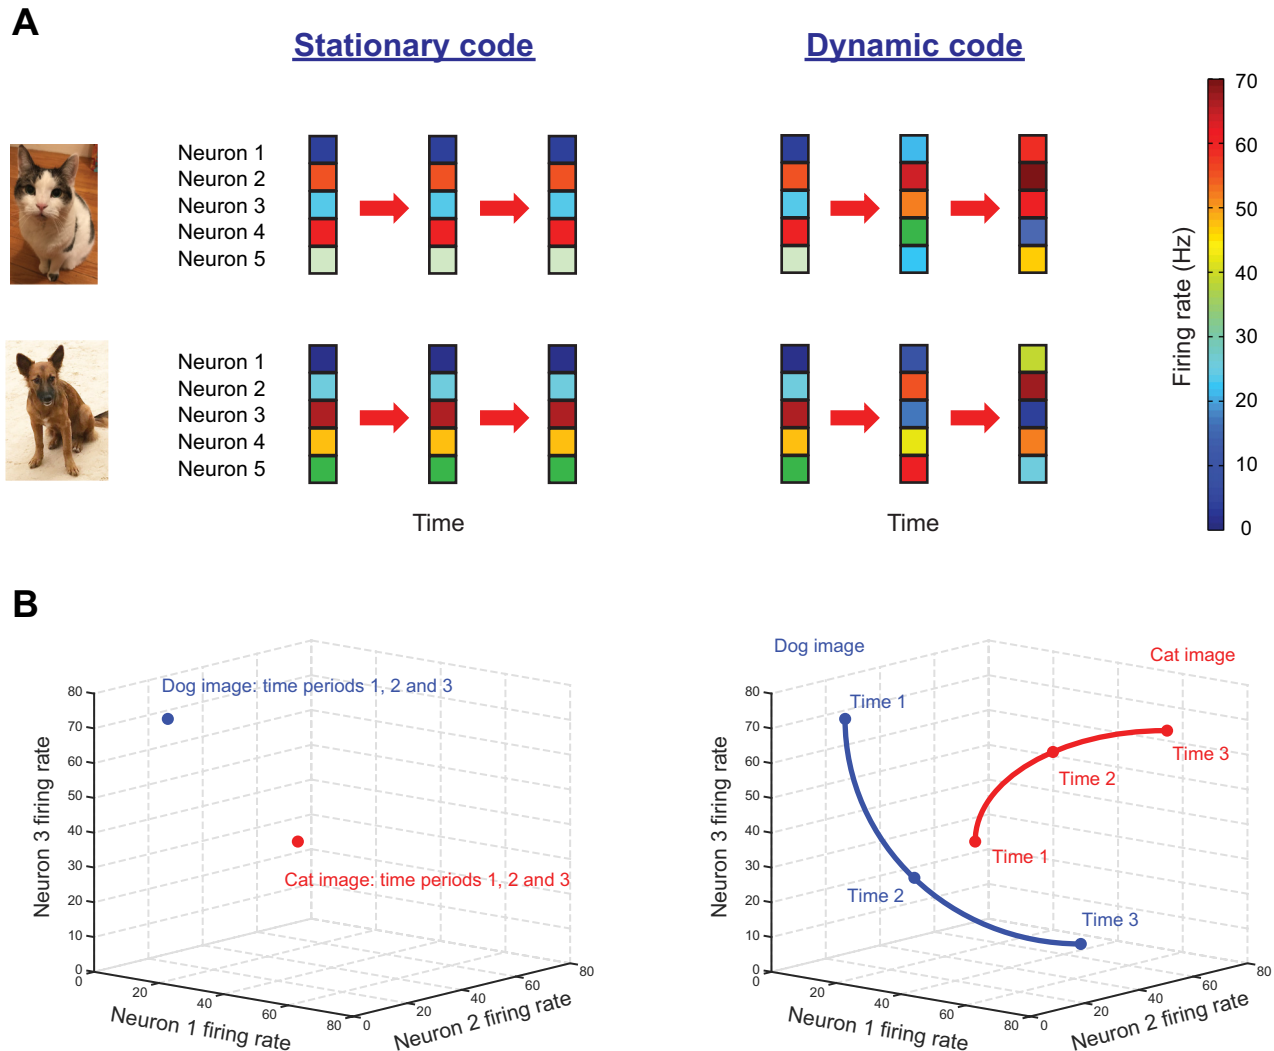
Fig. 2. Illustrations of stationary and dynamic codes. *A*: illustration of how a population vector of 5 neurons responds to an image of a cat (*top* plots) or a dog (*bottom* plots). Each neuron is shown as a square, and the color of the square indicates the firing rate of the neuron. For the stationary codes (shown on the *left*), the same pattern of neural activity across the 5 neurons is seen at all points in time, and these patterns differ depending on whether an image of a cat or a dog was shown. For the dynamic code (shown on the *right*), a different pattern of neural activity is present at each point in time, and the sequence of patterns also differs depending on whether an image of a cat or a dog was shown. *B*: geometric illustration of how a population of 3 neurons responds to the images of a cat (red) and a dog (blue). Each axis shows the firing rates of a neuron, and thus the population response is a point in 3-dimensional space. For a stationary code (shown on the *left*), the response to each image is a single point that does not change with time. For a dynamic code (shown on the *right*), the population response creates a trajectory in 3-dimensional space. Because these trajectories do not cross, it is possible to distinguish between cat and dog images at all points in time. It should be noted that in the literature, the terms stationary code, "stable code," and "static code" are used interchangeably to refer to the phenomenon illustrated on the *left* plots in this figure.

Beyond the analyses of neural spiking activity, cross-temporal decoding analyses of electroencephalography (EEG) and magnetoencephalography (MEG) data also revealed dynamic activity in those data, although changing patterns of EEG/MEG activity could be due to information traveling between brain regions rather than a dynamic code within a population of neurons (Carlson et al. 2011, 2013; Cichy et al. 2014; Isik et al. 2014; King and Dehaene 2014). Additionally, analyses of local field potentials in monkeys have revealed γ-burst during memory delay periods, which appear to be associated with dynamic codes (Lundqvist et al. 2016). Finally, Harvey et al. (2012) used two-photon imaging of mouse posterior parietal neurons and observed that most neurons were selective for relatively short periods of time and showed that neurons with different windows of selectivity were intermingled throughout this brain region.

It is important to note that the difference between studies that found stationary and dynamic codes is not due to different data analysis methods, since several studies using the temporal cross-decoding analyses have found stationary codes as well (see Fig. 4 and King and Dehaene 2014). Currently, it is not clear why some studies find dynamic code and others find stationary codes, although there seems to be a general trend that stationary codes are present when simple stimulus attributes need to be remembered (Meyers et al. 2017; Murray et al. 2017; Spaak et al. 2017; Zhang et al. 2011), whereas dynamic codes are present when more complex stimulus transformations are needed to solve a task (Crowe et al. 2010; Meyers et al. 2008, 2012; Stokes et al. 2013). Additionally, it appears that there might be more persistent sustained selectivity in the anterior-dorsal region of the prefrontal cortex, whereas there is
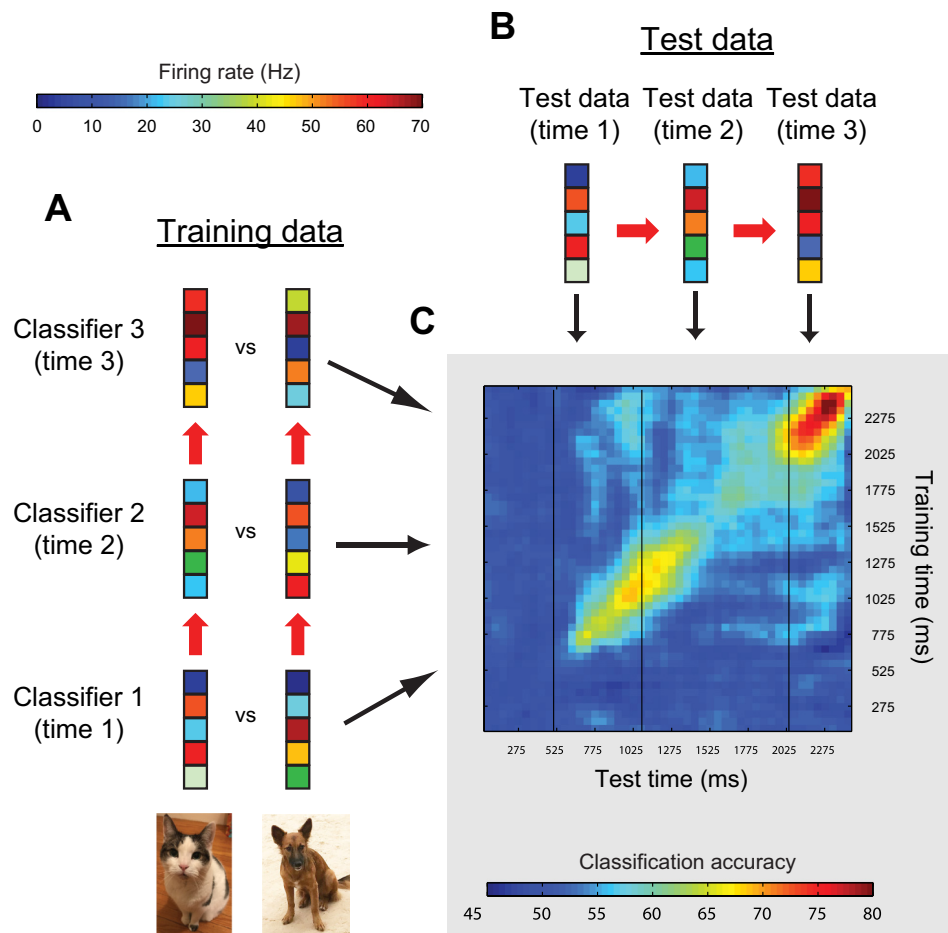
Fig. 3. Illustration of how a "cross-temporal decoding" analysis can be used to examine whether a stationary or dynamic code is present. *A*: for this analysis, a different classifier is trained to discriminate between images of dogs and cats at each point in time (indicated in the *y*-axis). *B*: these classifiers then need to predict whether an image of a dog or cat is present at each point in time on a different test set of data (indicated on the *x*-axis). *C*: cross-temporal decoding plot where the classification accuracy in indicated by the color bar in the bottom of the gray region. In this example, a high classification accuracy is only seen when the classifier is trained and tested at the same point in time, as can be seen by the fact that there is a colorful diagonal in the plot. This means that a dynamic code is present since a classifier trained at 1 point in time does not have a good performance when tested at a different point in time. If a stationary code was present, then one would expect to see a solid yellow/red colored square indicating good classification performance regardless of the time when the classifier was trained (see Fig. 4*B*). Black, vertical lines correspond to sample stimulus onset, sample stimulus offset, and match stimulus onset. [Results in this plot are adapted from Meyers et al. (2008).]

a more dynamic representation in the more posterior-dorsal regions (Leavitt et al. 2017; Meyers et al. 2017). If this is the case, then perhaps the stationary codes in the anterior-dorsal PFC are acting more like short-term memory and the dynamic codes in the posterior dorsolateral PFC region are more involved in manipulating information.

IS DYNAMIC ACTIVITY SUFFICIENT FOR MAINTAINING INFORMATION IN WORKING MEMORY?

The previous sections showed that both dynamic and stationary coding are robust phenomena that have been reported in many studies of working memory. However, definitively understanding the degree to which dynamic and stationary codes contribute to working memory is challenging. Nevertheless, it is possible to do analyses that can assess whether neurons that have dynamic coding properties are sufficient for maintaining information in working memory and whether such dynamic codes are "biologically plausible" in the sense that the information contained in dynamic codes could be utilized by downstream neurons to contribute to behavior. If it turns out that dynamic codes do not contain sufficient information for behavior or that information cannot be extracted from dynamic codes in a biologically plausible way, then clearly such coding would not be a viable candidate for the neural basis of working memory.

To assess whether dynamic neurons contain sufficient information to support working memory, Baeg et al. (2003) did an

analysis where they found all selective neurons that showed continuous activity in the delay period (where continuous activity was defined as neurons that had higher average firing rates in 4 bins in the delay period compared with other time periods). They then did a decoding analysis with these continuously active neurons removed and showed that the decoding accuracy level was similar using only the remaining dynamic neurons. Similarly, Meyers et al. (2008) did a decoding analysis where they eliminated all selective neurons based on the training data at 1 time point (which I will call t_{elim} here). This removed all neurons with persistent sustained selectivity from the population and removed all neurons with short dynamic periods of selectivity around t_{elim}. As expected, when the classifier was trained and tested with data from t_{elim}, the classification accuracy was around chance. However, when the classifier was trained and tested at other points in time with the same neurons from t_{elim} removed, the classification accuracy was once again high because there were different neurons that had short dynamic periods of selectivity at other points in time that the classifier could use (Fig. 5). These results show that neurons with shorter periods of selectivity seem to be sufficient to support working memory.

However, even if the dynamic neurons are sufficient for maintaining information in working memory, another question that arises is how dynamic codes could be readout by downstream neurons. In particular, artificial neural network models view neurons as being similar to generalized linear classifiers,

Fig. 4. Examples of dynamic and stationary codes in the parietal and prefrontal cortex. *A*: cross-temporal decoding plot for decoding whether 2 shapes matched during the delay period in a delayed match-to-sample task. Monkeys needed to saccade to a green target if the stimuli matched and a blue target if the stimuli did not match. Black, vertical lines indicate the time when the 1st, 2nd, and choice targets were shown. Chance decoding is 1/2 in this study. In this study, a highly dynamic code was present as indicated by the strong diagonal band in the plot (results adapted from Meyers et al. 2012). *B*: results from a temporal cross-decoding plot for decoding a location where a small square was shown during the 1st delay period in a delayed match-to-sample task. Monkeys needed to release a lever if a 2nd stimulus was shown at the same location. Since there were 4 locations where the stimulus could be shown, chance decoding is 1/4. In this study, a highly stationary code was present as indicated by a square region of high decoding accuracies at all time points. [Adapted from Meyers EM, Qi XL, Constantinidis C. Incorporation of new information into prefrontal cortical activity after learning working memory tasks. *Proc Natl Acad Sci USA* 109: 4651–4656, 2012, with permission.]
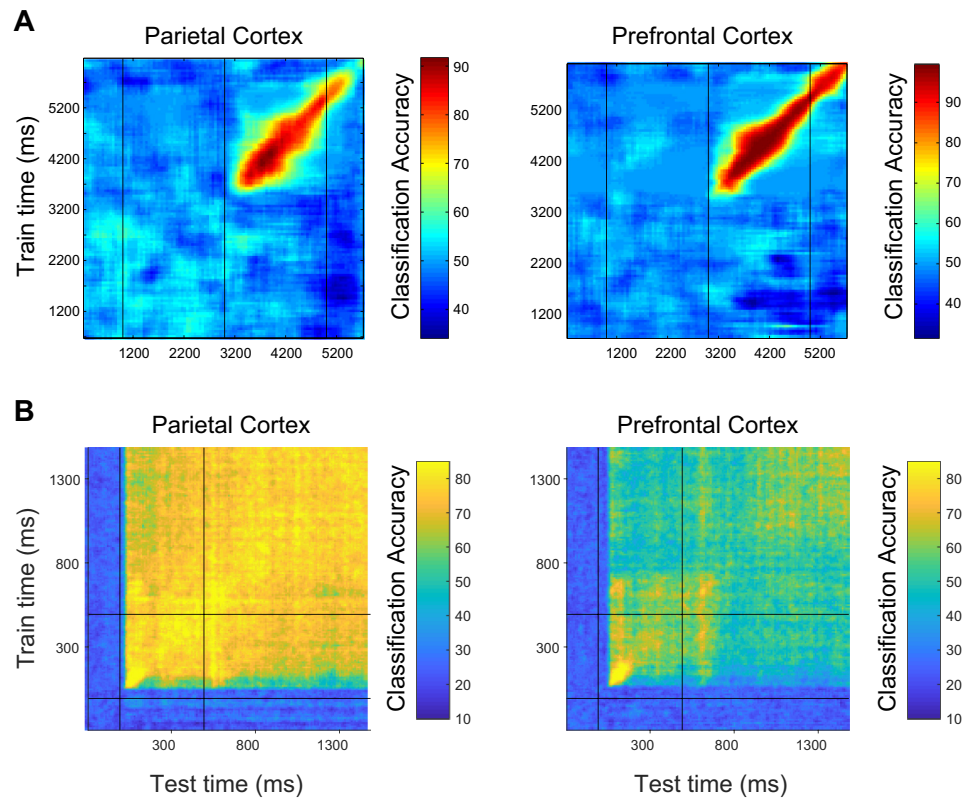
where the linear weights the classifier uses are analogous to the synaptic connections of a neuron (Fig. 6*A*). From this perspective, having downstream neurons readout information from a dynamic coding seems biologically problematic since it would imply that the synaptic weights would have to change continuously in time in a consistent manner on every single trial for the information to be continuously available to downstream brain regions. Whereas some authors have proposed that such dynamic synaptic changes might indeed be occurring (Stokes 2015), others have instead tried to assess whether it is possible to extract information from dynamic codes using a potentially more biologically plausible "fixed set" of synaptic weights (Meyers et al. 2009; Murray et al. 2017). For example, Meyers et al. (2009) examined a population of neurons that showed dynamic coding properties and tried to decode information from this population using a classifier that had one single fixed set of weights for all time points. More specifically, a classifier was trained to learn classification weights using neural activity that was averaged over the whole 2.5 s of an experimental trial. The classifier was tested (on a different set of trials) using sliding 150-ms bins. With the use of this approach, it was found that most of the information from the population could be extracted using a fixed set of weights even though many neurons were only selective for short periods of time (Fig. 6*B*). The reason why most of the information could be extracted using a fixed set of weights was because including activity outside of the window of time when a neuron was selective did not add enough noise to disrupt the information that was contained in the rest of the population. Thus most of the information from the population could still be extracted. Additionally, these results imply that even if downstream neurons integrate information over relatively long time scales of ~2 s,

most of the information from the upstream dynamic population is still available.

HOW COULD DYNAMIC CODES CONTRIBUTE TO WORKING MEMORY AND BEHAVIOR?

Taken together, these results show that dynamic codes have sufficient information to support working memory and that information contained in dynamic codes can be extracted by downstream neurons in a biologically plausible way. However, from a computational point of view, a question that arises is why would the brain use a dynamic code to hold information rather than a stationary code? In particular, it makes intuitive sense that persistent sustained selectivity would be the neural basis of working memory since encoding items in working memory would merely consist of retaining the original input activity and it would be trivial to retrieve the memory trace at any time point (Wimmer et al. 2014), whereas having dynamic codes seems to be more complicated. Indeed, it has been suggested that the real basis of working memory is, in fact, persistent activity and that these dynamic coding properties might be there to support the maintenance of persistent activity or purely epiphenomenal (Riley and Constantinidis 2016), although other researchers have taken the opposite perspective and have proposed that dynamics are the main basis of working memory (Lundqvist et al. 2016). Although it is currently hard to know whether either of these perspectives is correct, there are several theories that can explain why dynamic coding is useful. These theories can be broadly divided into two categories: namely, *1*) that dynamic codes are used to encode information about time or *2*) that dynamic codes are part of a sequence of processing steps needed to solve particular tasks.
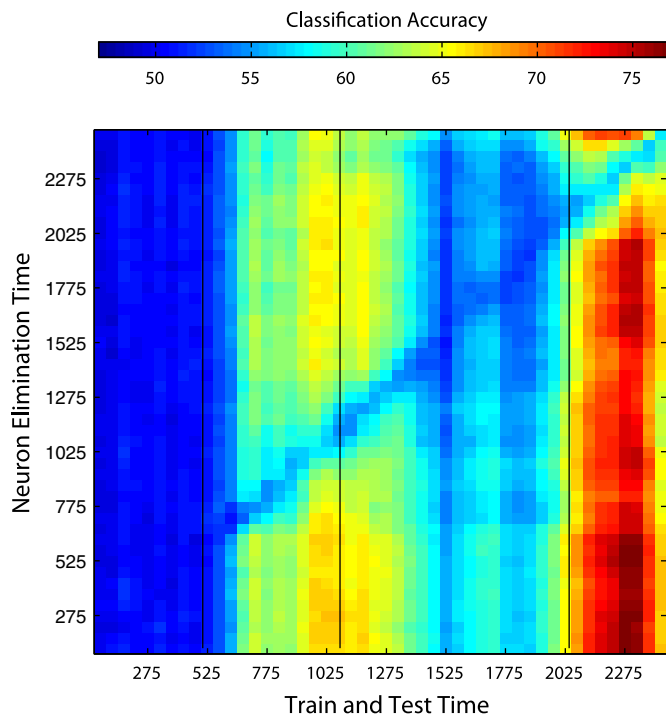
Classification Accuracy

Fig. 5. Results from an analysis examining whether different neurons are selective at different points in time. In this analysis, for each point in time, the 64 most selective neurons found in the training set were eliminated; the time point when the neurons were eliminated is shown in the y-axis. Classifier was then trained and tested with the remaining 192 neurons. As can be seen, the decoding accuracy had a large drop at the time point when the best neurons were eliminated, but then the decoding accuracy recovered at other points in time. Because the decoding accuracy recovered, this indicates that different neurons were selective at different points in time; otherwise, eliminating the best neurons at 1 point in time would cause poor performance at all other points in time. [Results are adapted from Meyers et al. (2008).]

The first class of theories, that dynamic coding is related to how the brain keeps track of time, has been put forward by several research groups. For example, researchers studying the hippocampus have proposed that these dynamic neurons are "time cells" that are used for recalling temporal memories (Eichenbaum 2014; Pastalkova et al. 2008). Similarly, researchers examining a task where monkeys needed to produce responses at particular time intervals found neurons in the medial frontal cortex that were tuned to the timing of the task, again suggesting that the dynamics are involved in keeping track of time (Wang et al. 2018). If this is the case, then the dynamic coding seen in working memory could merely be a result of information about time that is being multiplexed on top of memory signals containing information about other variables. Many studies have shown that PFC activity simultaneously contains information about multiple behaviorally relevant variables, and it has been suggested that the complex coding seen in PFC might be a useful feature for separating all possible stimulus configurations (Rigotti et al. 2013), which could include information about time as well.

If dynamic population codes are a mechanism for keeping track of time, then it seems that (at least in the PFC) information about time is contained relative to particular events of interest rather than on an absolute scale (Brody et al. 2003a; Wang et al. 2018). In particular, Brody et al. (2003a), analyzed dynamic activity in PFC in an experiment where they com-

pared neural activity in a 3-s delay period to neural activity in a 6-s delay period (the experiment was done in a blocked design so that the monkey could quickly adapt its behavior to the new delay period length). Their results showed that the dynamic neural activity appears to evolve slower when the delay period was longer, with neural activity in the 6-s delay period being well-modeled by a stretching of the neural activity seen in the 3-s delay period. Thus, from a population perspective, the same dynamic codes would occur at the same relative position in the delay period regardless of the duration of the delay period, and so information about time would be relative to the start and end of the delay period.

The second class of theories, that dynamic codes are related to different steps needed to solve a task, has been be formulated in several related ways. One formulation suggests that the temporal dynamics are actually due to neurons coding different information at different points in time; however, because information about several variables are often confounded in a typical experiment, it could (falsely) appear that information about one single variable is being coded dynamically. For example, imagine a delayed match-to-sample task in which a monkey needs to remember a location of a cue stimulus and then saccade to a particular target if the second stimulus is at the same location. During the memory delay period in this task, a population of neurons could initially code information about the location of the first cue stimulus. However, then, later in the delay, the population could contain a signal related to triggering a motor response if the second stimulus is at the matching location. Since information about the first stimulus and triggering a motor response in anticipation of a second stimulus are confounded, then it might look like information about the first stimulus is stored in memory by a dynamic code, but really the dynamics are due to the fact that different information is being coded at different times in a trial. This interpretation is also consistent with the fact that more dynamic coding is seen in experiments that involve more complex transformation of stimulus attributes, which could be due to the neural population representing information in intermediate computational steps (Crowe et al. 2010; Meyers et al. 2008, 2012; Stokes et al. 2013). Such sequential steps of processing are also similar to the description of the PFC as an area that is responsible for integrating information over time (Fuster 2001).

A related formulation is the notion that neural processing is best described as a dynamic system that goes through a sequence of hidden states to solve particular tasks (Shenoy et al. 2013). In particular, dynamic representations have been seen in dorsal premotor cortex when monkeys are preparing to make motor movements, and monkeys' reaction times can be predicted based on position and rate of change of dynamic coding activity, which shows there is a strong association between neural dynamics and behavior (Afshar et al. 2011). This dynamic system explanation is similar to the "confounding" explanation described in the previous paragraph in the sense that both explanations describe states that are hidden from the experimentalists, although the dynamic system perspective might not try to assign meaning to different internal states that could be continuously changing.

The concept that dynamic coding is related to different computational steps also makes sense from the perspective that it would be impossible to solve a particular task if the same pattern of neural activity that initially encoded the stimulus was
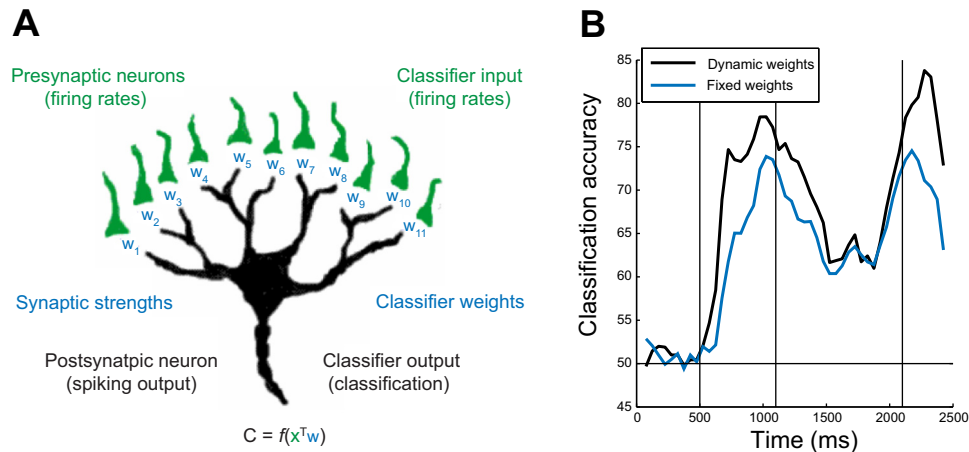
A

Presynaptic neurons
(firing rates)

Classifier input
(firing rates)

w_5 w_6 w_7 w_8
w_4
w_3 w_9 w_{10}
w_2 w_{11}
w_1

Synaptic strengths   Classifier weights

Postsynatpic neuron
(spiking output)

Classifier output
(classification)

$$C = f(x^T w)$$

B

Fig. 6. Results showing that it is possible to extract most of the information from a dynamic code using a fixed set of classification weights. *A*: illustration of the correspondence between a linear classifier and a simple model of a neuron. Here, the presynaptic activity corresponds to the input of the classifier (green), the synaptic strengths correspond to the weights of the classifier (w; blue), and the spiking output of the neuron corresponds to the prediction of the classifier (black), as indicated by the equation $C = f(x^T w)$, where f is a nonlinear function applied to the linear combination of synaptic strengths w and neural firing rate input x. *B*: results showing the classification accuracy for a dynamic code when a different set of weights was used at each point in time (black trace) or the same set of weights was used at all points in time (blue trace). Results plotted in the black trace using a different classifier at each time point are the same results as are plotted on the diagonal of Fig. 3*C*. Results in the blue trace are from training a classifier using 1 large time bin that has the average activity over the whole trial, and then the classifier was tested on smaller sliding bins. As can be seen, the decoding accuracy is slightly lower when a fixed set of weights is used at all time points (blue trace); however, this classifier still captures most of the information in the population despite the fact that a dynamic code is present. The reason this works is because even if neurons are selective for short periods of time, the activity outside of these short windows of selectivity does not degrade the classification performance at other points in time. Thus it is plausible that a downstream neuron with a fixed set of weights could extract most of the information from an upstream dynamic population code. [Results are adapted from a poster presented at the 2009 *COSYNE* conference (abstract published as Meyers et al. 2009).]

also used for maintaining the stimulus in working memory. For example, consider again the delayed match-to-sample task. In this task, if the same pattern of neural activity was used to represent both the first and second stimuli, then information about the first stimulus would be overwritten when the second stimulus was shown, which would make it impossible to tell whether the two stimuli matched. Conversely, if dynamic coding was used, then it would be possible to compare the same stimulus across time because different patterns of activity would represent the first and second occurrences of the stimulus. Thus a dynamic coding scheme seems necessary, to some degree, to solve particular tasks.

Finally, it should be noted that it is also possible that the dynamic coding might be epiphenomenal and is not enabling an animal to perform any additional behaviors. As described above, dynamic population activity can be readout by downstream neurons that integrate over the full delay period without much loss of information (Meyers et al. 2009), so downstream neurons might indeed be trying to extract constant information from the population in the face of these dynamics rather than having these dynamic codes as a useful feature that contributes meaningfully to particular behaviors, such as keeping track of time (Murray et al. 2017).

CONCLUSIONS

In this paper, I have reviewed literature showing that both persistent sustained activity and dynamic population coding are associated with working memory. Furthermore, I have described studies that have shown that the amount of information stored in dynamic codes appears to be sufficient to solve working memory tasks and that this information could be extracted from a neural population by downstream neurons using biologically plausible mechanisms. Finally, I explored

possible explanations for why dynamic codes might be present and how they could potentially contribute to solving tasks.

Throughout this paper, I have described neural coding as being either persistent or dynamic. However, it should be pointed out that on an individual neuron level there often seems to be a continuum of time scale that neurons encode information (Harvey et al. 2012; Meyers et al. 2008, 2012; Zaksas and Pasternak 2006), although there might also be systematic differences in time scales as well between different brain regions or even different animals (Leavitt et al. 2017; Mendoza-Halliday et al. 2014; Meyers et al. 2017; Murray et al. 2014; Runyan et al. 2017). Also, it should be mentioned that estimating the time scale that a neuron maintains information could be imprecise in many studies because information might be present at different points in time on different trials. Thus, when data are analyzed across trials, it could appear that information is maintained for longer periods of time than it really is on any single individual trial. By analyzing larger populations of neurons that are recorded simultaneously, it might be possible to track information content changes within a single trial, which could better help estimate how long neurons typically maintain information for.

In summary, although many questions remain about how persistent activity and dynamic coding contribute to working memory, it is clear that both types of coding are commonly seen phenomena that carry large amounts of information useful for solving tasks. Although much work remains, the facts that recording technologies continue to improve and that data analysis toolboxes exist that make it relatively easy to assess dynamics in neural spiking and EEG/MEG activity (Gramfort et al. 2013; Meyers 2013; Tadel et al. 2011) should lead to a more concrete description of the neural algorithms that underlie particular behaviors in the not too distant future.

ACKNOWLEDGMENTS

GRANTS

DISCLOSURES

No conflicts of interest, financial or otherwise, are declared by the author.

AUTHOR CONTRIBUTIONS

E.M.M. prepared figures and drafted, edited, revised, and approved the final version of the manuscript.

REFERENCES

Aben B, Stapert S, Blokland A. About the distinction between working memory and short-term memory. *Front Psychol* 3: 301, 2012. doi:10.3389/fpsyg.2012.00301.

Afshar A, Santhanam G, Yu BM, Ryu SI, Sahani M, Shenoy KV. Single-trial neural correlates of arm movement preparation. *Neuron* 71: 555–564, 2011. doi:10.1016/j.neuron.2011.05.047.

Baddeley A. Working memory: theories, models, and controversies. *Annu Rev Psychol* 63: 1–29, 2012. doi:10.1146/annurev-psych-120710-100422.

Baeg EH, Kim YB, Huh K, Mook-Jung I, Kim HT, Jung MW. Dynamics of population code for working memory in the prefrontal cortex. *Neuron* 40: 177–188, 2003. doi:10.1016/S0896-6273(03)00597-X.

Barak O, Tsodyks M, Romo R. Neuronal population coding of parametric working memory. *J Neurosci* 30: 9424–9430, 2010. doi:10.1523/JNEURO-SCI.1875-10.2010.

Batuev AS, Pirogov AA, Orlov AA, Sheafer VI. Cortical mechanisms of goal-directed motor acts in the rhesus monkey. *Acta Neurobiol Exp (Wars)* 40: 27–49, 1980.

Brody CD, Hernández A, Zainos A, Romo R. Timing and neural encoding of somatosensory parametric working memory in macaque prefrontal cortex. *Cereb Cortex* 13: 1196–1207, 2003a. doi:10.1093/cercor/bhg100.

Brody CD, Romo R, Kepecs A. Basic mechanisms for graded persistent activity: discrete attractors, continuous attractors, and dynamic representations. *Curr Opin Neurobiol* 13: 204–211, 2003b. doi:10.1016/S0959-4388(03)00050-3.

Brunel N. Dynamics and plasticity of stimulus-selective persistent activity in cortical network models. *Cereb Cortex* 13: 1151–1161, 2003. doi:10.1093/cercor/bhg096.

Carlson T, Tovar DA, Alink A, Kriegeskorte N. Representational dynamics of object vision: the first 1000 ms. *J Vis* 13: 1, 2013. doi:10.1167/13.10.1.

Carlson TA, Hogendoorn H, Kanai R, Mesik J, Turret J. High temporal resolution decoding of object position and category. *J Vis* 11: 9, 2011. doi:10.1167/11.10.9.

Chafee MV, Goldman-Rakic PS. Matching patterns of activity in primate prefrontal area 8a and parietal area 7ip neurons during a spatial working memory task. *J Neurophysiol* 79: 2919–2940, 1998. doi:10.1152/jn.1998.79.6.2919.

Cichy RM, Pantazis D, Oliva A. Resolving human object recognition in space and time. *Nat Neurosci* 17: 455–462, 2014. doi:10.1038/nn.3635.

Constantinidis C, Wang XJ. A neural circuit basis for spatial working memory. *Neuroscientist* 10: 553–565, 2004. doi:10.1177/1073858404268742.

Cowan N. What are the differences between long-term, short-term, and working memory? *Prog Brain Res* 169: 323–338, 2008. doi:10.1016/S0079-6123(07)00020-9.

Crowe DA, Averbeck BB, Chafee MV. Rapid sequences of population activity patterns dynamically encode task-critical spatial information in parietal cortex. *J Neurosci* 30: 11640–11653, 2010. doi:10.1523/JNEURO-SCI.0954-10.2010.

D'Esposito M, Postle BR. The cognitive neuroscience of working memory. *Annu Rev Psychol* 66: 115–142, 2015. doi:10.1146/annurev-psych-010814-015031.

Durstewitz D, Seamans JK, Sejnowski TJ. Neurocomputational models of working memory. *Nat Neurosci* 3, *Suppl*: 1184–1191, 2000. doi:10.1038/81460.

Eichenbaum H. Time cells in the hippocampus: a new dimension for mapping memories. *Nat Rev Neurosci* 15: 732–744, 2014. doi:10.1038/nrn3827.

Funahashi S, Chafee MV, Goldman-Rakic PS. Prefrontal neuronal activity in rhesus monkeys performing a delayed anti-saccade task. *Nature* 365: 753–756, 1993. doi:10.1038/365753a0.

Fuster JM. The prefrontal cortex–an update: time is of the essence. *Neuron* 30: 319–333, 2001. doi:10.1016/S0896-6273(01)00285-9.

Fuster JM. Unit activity in prefrontal cortex during delayed-response performance: neuronal correlates of transient memory. *J Neurophysiol* 36: 61–78, 1973. doi:10.1152/jn.1973.36.1.61.

Fuster JM, Alexander GE. Neuron activity related to short-term memory. *Science* 173: 652–654, 1971. doi:10.1126/science.173.3997.652.

Goldman-Rakic PS. Regional and cellular fractionation of working memory. *Proc Natl Acad Sci USA* 93: 13473–13480, 1996. doi:10.1073/pnas.93.24.13473.

Gramfort A, Luessi M, Larson E, Engemann DA, Strohmeier D, Brodbeck C, Goj R, Jas M, Brooks T, Parkkonen L, Hämäläinen M. MEG and EEG data analysis with MNE-Python. *Front Neurosci* 7: 267, 2013. doi:10.3389/fnins.2013.00267.

Harvey CD, Coen P, Tank DW. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* 484: 62–68, 2012. doi:10.1038/nature10918.

Isik L, Meyers EM, Leibo JZ, Poggio T. The dynamics of invariant object recognition in the human visual system. *J Neurophysiol* 111: 91–102, 2014. doi:10.1152/jn.00394.2013.

King JR, Dehaene S. Characterizing the dynamics of mental representations: the temporal generalization method. *Trends Cogn Sci* 18: 203–210, 2014. doi:10.1016/j.tics.2014.01.002.

Kubota K, Niki H. Prefrontal cortical unit activity and delayed alternation performance in monkeys. *J Neurophysiol* 34: 337–347, 1971. doi:10.1152/jn.1971.34.3.337.

Leavitt ML, Mendoza-Halliday D, Martinez-Trujillo JC. Sustained activity encoding working memories: not fully distributed. *Trends Neurosci* 40: 328–346, 2017. doi:10.1016/j.tins.2017.04.004.

Lundqvist M, Rose J, Herman P, Brincat SL, Buschman TJ, Miller EK. Gamma and beta bursts underlie working memory. *Neuron* 90: 152–164, 2016. doi:10.1016/j.neuron.2016.02.028.

Major G, Tank D. Persistent neural activity: prevalence and mechanisms. *Curr Opin Neurobiol* 14: 675–684, 2004. doi:10.1016/j.conb.2004.10.017.

Mendoza-Halliday D, Torres S, Martinez-Trujillo JC. Sharp emergence of feature-selective sustained activity along the dorsal visual pathway. *Nat Neurosci* 17: 1255–1262, 2014. doi:10.1038/nn.3785.

Meyers E, Freedman D, Kreiman G, Miller E, Poggio T. Decoding dynamic patterns of neural activity using a 'biologically plausible' fixed set of weights (Abstract). Computational and Systems Neuroscience 2009. Salt Lake City, UT: February 26 to March 3, 2009. doi:10.3389/conf.neuro.06.2009.03.354.

Meyers E, Riley M, Qi XL, Constantinidis C. Differences in dynamic and static coding within different subdivision of the prefrontal cortex (Abstract). Society for Neuroscience's Annual Meeting – SfN 2017. Washington, DC, 2017.

Meyers EM. The neural decoding toolbox. *Front Neuroinform* 7: 8, 2013. doi:10.3389/fninf.2013.00008.

Meyers EM, Freedman DJ, Kreiman G, Miller EK, Poggio T. Dynamic population coding of category information in inferior temporal and prefrontal cortex. *J Neurophysiol* 100: 1407–1419, 2008. doi:10.1152/jn.90248.2008.

Meyers EM, Qi XL, Constantinidis C. Incorporation of new information into prefrontal cortical activity after learning working memory tasks. *Proc Natl Acad Sci USA* 109: 4651–4656, 2012. doi:10.1073/pnas.1201022109.

Miller EK, Erickson CA, Desimone R. Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *J Neurosci* 16: 5154–5167, 1996. doi:10.1523/JNEUROSCI.16-16-05154.1996.

Mongillo G, Barak O, Tsodyks M. Synaptic theory of working memory. *Science* 319: 1543–1546, 2008. doi:10.1126/science.1150769.

Murray JD, Bernacchia A, Freedman DJ, Romo R, Wallis JD, Cai X, Padoa-Schioppa C, Pasternak T, Seo H, Lee D, Wang XJ. A hierarchy of intrinsic timescales across primate cortex. *Nat Neurosci* 17: 1661–1663, 2014. doi:10.1038/nn.3862.

Murray JD, Bernacchia A, Roy NA, Constantinidis C, Romo R, Wang XJ. Stable population coding for working memory coexists with heterogeneous

neural dynamics in prefrontal cortex. *Proc Natl Acad Sci USA* 114: 394–399, 2017. doi:10.1073/pnas.1619449114.

Nee DE, D'Esposito M. Working memory. In: *Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience*. Hoboken, NJ: John Wiley & Sons, 2018, p. 1–26. doi:10.1002/9781119170174.epcn112.

Nikolić D, Häusler S, Singer W, Maass W. Distributed fading memory for stimulus properties in the primary visual cortex. *PLoS Biol* 7: e1000260, 2009. doi:10.1371/journal.pbio.1000260.

Pastalkova E, Itskov V, Amarasingham A, Buzsáki G. Internally generated cell assembly sequences in the rat hippocampus. *Science* 321: 1322–1327, 2008. doi:10.1126/science.1159775.

Rigotti M, Barak O, Warden MR, Wang XJ, Daw ND, Miller EK, Fusi S. The importance of mixed selectivity in complex cognitive tasks. *Nature* 497: 585–590, 2013. doi:10.1038/nature12160.

Riley MR, Constantinidis C. Role of prefrontal persistent activity in working memory. *Front Syst Neurosci* 9: 181, 2016. doi:10.3389/fnsys.2015.00181.

Romo R, Brody CD, Hernández A, Lemus L. Neuronal correlates of parametric working memory in the prefrontal cortex. *Nature* 399: 470–473, 1999. doi:10.1038/20939.

Runyan CA, Piasini E, Panzeri S, Harvey CD. Distinct timescales of population coding across cortex. *Nature* 548: 92–96, 2017. doi:10.1038/nature23020.

Shafi M, Zhou Y, Quintana J, Chow C, Fuster J, Bodner M. Variability in neuronal activity in primate cortex during working memory tasks. *Neuroscience* 146: 1082–1108, 2007. doi:10.1016/j.neuroscience.2006.12.072.

Shenoy KV, Sahani M, Churchland MM. Cortical control of arm movements: a dynamical systems perspective. *Annu Rev Neurosci* 36: 337–359, 2013. doi:10.1146/annurev-neuro-062111-150509.

Spaak E, Watanabe K, Funahashi S, Stokes MG. Stable and dynamic coding for working memory in primate prefrontal cortex. *J Neurosci* 37: 6503–6516, 2017. doi:10.1523/JNEUROSCI.3364-16.2017.

Sreenivasan KK, Curtis CE, D'Esposito M. Revisiting the role of persistent neural activity during working memory. *Trends Cogn Sci* 18: 82–89, 2014. doi:10.1016/j.tics.2013.12.001.

Stokes MG. 'Activity-silent' working memory in prefrontal cortex: a dynamic coding framework. *Trends Cogn Sci* 19: 394–405, 2015. doi:10.1016/j.tics.2015.05.004.

Stokes MG, Kusunoki M, Sigala N, Nili H, Gaffan D, Duncan J. Dynamic coding for cognitive control in prefrontal cortex. *Neuron* 78: 364–375, 2013. doi:10.1016/j.neuron.2013.01.039.

Tadel F, Baillet S, Mosher JC, Pantazis D, Leahy RM. Brainstorm: a user-friendly application for MEG/EEG analysis. *Comput Intell Neurosci* 2011: 879716, 2011. doi:10.1155/2011/879716.

Wang J, Narain D, Hosseini EA, Jazayeri M. Flexible timing by temporal scaling of cortical responses. *Nat Neurosci* 21: 102–110, 2018. doi:10.1038/s41593-017-0028-6.

Wang XJ. Synaptic reverberation underlying mnemonic persistent activity. *Trends Neurosci* 24: 455–463, 2001. doi:10.1016/S0166-2236(00)01868-3.

Warden MR, Miller EK. The representation of multiple objects in prefrontal neuronal delay activity. *Cereb Cortex* 17, *Suppl* 1: i41–i50, 2007. doi:10.1093/cercor/bhm070.

Wimmer K, Nykamp DQ, Constantinidis C, Compte A. Bump attractor dynamics in prefrontal cortex explains behavioral precision in spatial working memory. *Nat Neurosci* 17: 431–439, 2014. doi:10.1038/nn.3645.

Zaksas D, Pasternak T. Directional signals in the prefrontal cortex and in area MT during a working memory for visual motion task. *J Neurosci* 26: 11726–11742, 2006. doi:10.1523/JNEUROSCI.3420-06.2006.

Zhang Y, Meyers EM, Bichot NP, Serre T, Poggio TA, Desimone R. Object decoding with attention in inferior temporal cortex. *Proc Natl Acad Sci USA* 108: 8850–8855, 2011. doi:10.1073/pnas.1100999108.

Zylberberg J, Strowbridge BW. Mechanisms of persistent activity in cortical circuits: possible neural substrates for working memory. *Annu Rev Neurosci* 40: 603–627, 2017. doi:10.1146/annurev-neuro-070815-014006.