



Published in final edited form as:

Annu Rev Neurosci. 2017 July 25; 40: 603–627. doi:10.1146/annurev-neuro-070815-014006.

Mechanisms of Persistent Activity in Cortical Circuits: Possible Neural Substrates for Working Memory

Joel Zylberberg^{1,2,3} and Ben W. Strowbridge^{4,5}

¹Department of Physiology and Biophysics, Center for Neuroscience, and Computational Bioscience Program, University of Colorado School of Medicine, Aurora, Colorado 80045

²Department of Applied Mathematics, University of Colorado, Boulder, Colorado 80309

³Learning in Machines and Brains Program, Canadian Institute for Advanced Research, Toronto, Ontario M5G 1Z8, Canada

⁴Department of Neurosciences, Case Western Reserve University School of Medicine, Cleveland, Ohio 44106

⁵Department of Physiology and Biophysics, Case Western Reserve University School of Medicine, Cleveland, Ohio 44106

Abstract

A commonly observed neural correlate of working memory is firing that persists after the triggering stimulus disappears. Substantial effort has been devoted to understanding the many potential mechanisms that may underlie memory-associated persistent activity. These rely either on the intrinsic properties of individual neurons or on the connectivity within neural circuits to maintain the persistent activity. Nevertheless, it remains unclear which mechanisms are at play in the many brain areas involved in working memory. Herein, we first summarize the palette of different mechanisms that can generate persistent activity. We then discuss recent work that asks which mechanisms underlie persistent activity in different brain areas. Finally, we discuss future studies that might tackle this question further. Our goal is to bridge between the communities of researchers who study either single-neuron biophysical, or neural circuit, mechanisms that can generate the persistent activity that underlies working memory.

Keywords

short-term memory; persistent activity; plateau potential; bistability; attractor network; feedback; synaptic transmission; neocortex

OVERVIEW

Working memory (WM), in which information about an external stimulus is stored by the brain for several seconds after that stimulus goes away, is a key cognitive function. For

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

example, to remember a phone number or infer causal relationships, one needs to recall the ordering of a sequence, and that ability hinges on WM. Decades of physiological studies have identified persistent firing as a neural correlate of WM. In the early, landmark studies in this field, monkeys were presented with a transient stimulus and trained to report that stimulus after a delay period in which no stimulus was present. Recordings in the monkeys' prefrontal cortices (PFCs) showed that some neurons showed sustained, elevated firing rates during the delay period and that the delay period firing rates depended systematically on the identity of the previously presented stimulus (Fuster 1973, Kubota et al. 1974). Subsequent studies similarly observed stimulus-tuned delay period activity in the PFC (Barak et al. 2010, Brody et al. 2003, Funahashi et al. 1989, Goldman-Rakic 1995, Miller et al. 1996, Rainer et al. 1998, Rao et al. 1997, Romo et al. 1999, Shafi et al. 2007) as well as in various other parts of the monkey brain: the temporal cortex (Bisley et al. 2004, Fuster & Jervey 1981, Miller et al. 1993, Miyashita & Chang 1988, Zaksas & Pasternak 2006), parietal cortex (Gnadt & Andersen 1988, Pesaran et al. 2002, Shafi et al. 2007), auditory cortex (Gottlieb et al. 1989), visual cortex (Super et al. 2001), somatosensory cortex (Zhou & Fuster 1996), presupplementary motor area (Vergara et al. 2016), and medial premotor cortex (Hernández et al. 2002). Recent work has expanded to rodents, in which genetic tools allow neural activities to be manipulated as well as recorded. That work has focused on rodent homologs of the primate brain areas involved in memory and movement planning, including the parietal cortex, PFC, anterior lateral motor cortex (ALM), frontal orienting fields (FOF), and superior colliculus (SC) (Harvey et al. 2012, Kopec et al. 2015, Li et al. 2016, Liu et al. 2014). In rodents, elevated delay period firing is observed as in the monkey, but relatively few cells are active for the entire delay period. Instead, most cells fire only in particular time windows. Those time windows tile the delay period, leading to observations of sequentially activated neurons (Harvey et al. 2012, MacDonald et al. 2011, Pastalkova et al. 2008), with the sets of sequentially activated neurons indicating the upcoming behavioral outcome (Harvey et al. 2012).

Supporting the role of elevated delay period firing in WM, this elevated firing is typically extinguished once the animal reports the remembered stimulus—at which point the information is no longer needed for the memory task (Funahashi et al. 1989, Fuster & Jervey 1981). Moreover, the delay period activity is preferentially tuned to those stimulus features that are needed for the subsequent behavior (Rainer et al. 1998, Rao et al. 1997), as opposed to those stimulus features that are less task-relevant. To causally test the role of delay period activity in WM, recent work has focused on single-trial analyses of neural activity and behavior and on the impact of optogenetic perturbations on the behavioral reports at the end of the delay period; as argued by Panzeri et al. (2017), these two techniques provide powerful tests of the causal role of neural activities in driving behavior. This work shows that, even over trials in which the animal had to remember the same presented stimulus, differences in the delay period activities are predictive of subsequent differences in the behavior (Kopec et al. 2015, Li et al. 2016, Vergara et al. 2016, Wimmer et al. 2014). Moreover, optogenetic perturbations to the neural activities during the delay period cause predictable changes in the subsequent behavior (Kopec et al. 2015, Li et al. 2016, Liu et al. 2014). Collectively, this body of work points to elevated neural firing during the delay period as an active process through which the brain stores information. The mechanisms underlying

this persistent activity are the focus of this review (see the sidebar titled Key Phrases Defined).

KEY PHRASES DEFINED

Persistent activity: neural firing that continues after the triggering stimulus goes away.

Bistability and multistability: forms of persistent activity in which a brief stimulus initiates a step-like change in firing rate. In bistable neurons, the change typically is from quiescence (no spontaneous spiking) to a quasi-constant firing rate. Multistable cells are capable of maintaining different stable firing rates.

Afterdepolarization: a response to a depolarizing stimulus in which the membrane potential does not return immediately to the prestimulus potential (it remains depolarized for some time, typically from several hundred milliseconds to several seconds). If large enough and if the resting membrane potential is close enough to the action potential threshold, the afterdepolarization can trigger persistent spiking.

Plateau potential: a classic term for an afterdepolarization response in which the membrane potential moves to a new (more depolarized) potential following a stimulus for several seconds before decaying back to the resting potential. Plateau potentials are often generated by noninactivating inward currents mediated by VGCCs.

Window current: current through an ion channel that occurs in the window between the membrane potential at which the ion channel activates but does not fully inactivate.

Feedforward connectivity: a pattern of connectivity in which the principal cells in one layer (or one brain region or subfield) do not form synaptic connections with each other. All lateral interactions among principal cells are mediated by interneurons.

Recurrent connectivity: connectivity between neurons that is not strictly feedforward because it includes connections between neurons of the same type (e.g., neocortical pyramidal cells that synapse onto other pyramidal cells). This connectivity allows for loops such as $A \rightarrow B \rightarrow A$ and $A \rightarrow B \rightarrow C \rightarrow A$ that are required for attractor network models.

Discrete attractor model: a neural circuit model in which a discrete set of stable neural firing rate patterns are generated. Upon stimulation, the network settles into one of these patterns: The specific pattern that arises depends on the stimulation applied to the circuit.

Continuous attractor model: similar to a discrete attractor model, but with a continuum of stable firing rate patterns.

Network eigenvalue: a number that quantifies the strength with which the recurrent circuitry in a neural network maintains a particular pattern of neural firing rates (that pattern is known as the eigenvector associated with the eigenvalue). This number quantifies the fraction of the amplitude of the neural firing rate pattern that is regenerated by the recurrent circuitry. Values less than 1 mean that the firing rate pattern decays over time, values of exactly 1 mean that the pattern is maintained at a constant level, and values above 1 mean that the amplitude of the pattern increases over time owing to the recurrent circuitry.

Fine-tuning: the observation that, for many neural circuit models of persistent activity based on recurrent connectivity, small deviations to the network connectivity (in either the pattern of connections or the connection weights) destroy the persistence.

Before we review the mechanisms that might underlie active WM, it is important to note that there are mechanisms through which information can be stored passively, without sustained neural spiking. For example, researchers have proposed that, when a stimulus evokes spiking activity in a neural circuit, short-term activity-dependent synaptic plasticity modifies that circuit so that it is primed to generate the same activity pattern again later (Mongillo et al. 2008, Santos et al. 2012). Consequently, the stimulus-evoked pattern is remembered by the pattern of synaptic modifications. Depending on the dynamics underlying the synaptic changes, this information can be maintained for relatively long periods of time (Benna & Fusi 2016, Fusi et al. 2005, Lahiri & Ganguli 2013). Models of this type are reviewed in more detail elsewhere (Barak & Tsodyks 2014), and so here we focus instead on the active mechanisms that involve persistent spiking.

Despite decades of study—see, for example, prior reviews on this subject (Brody et al. 2003, Chaudhuri & Fiete 2016, Compte 2006, Durstewitz et al. 2000, Hasselmo & Stern 2006)—we still lack a comprehensive understanding of which mechanisms generate the persistent activity that appears to underlie active WM. For example, is the persistent activity an intrinsic property of the neurons themselves (so that they keep spiking once they are appropriately activated; Egorov et al. 2002, Fraser & MacVicar 1996), does the persistence rely on the connectivity within the neural circuit to extend the representation beyond the timescale of activity in isolated neurons (Compte et al. 2000, Druckmann & Chklovskii 2012, Goldman 2009, Hopfield 1982, Lim & Goldman 2013, Seung 1996), or does the persistence depend instead on a combination of these mechanisms (Goldman et al. 2003, Koulakov et al. 2002)? Why have these questions not yet been answered? Several factors appear to be at play. First, multiple mechanistic models can give rise to similar observable neural activity patterns and (in some cases) to similar impacts of optogenetic perturbation on WM function. This degeneracy makes it hard to rule out specific mechanistic causes in favor of others. Moreover, several mechanisms may coexist within a brain area, and some mechanisms may exist only under certain neuromodulatory conditions, making it hard to say with certainty which ones are necessary, sufficient, or both. Finally, different brain areas might express these different mechanisms differentially, challenging attempts to apply the findings from one brain area to another.

To encourage progress on understanding the neurobiology of WM, we review here the palette of different cellular and network mechanisms that can give rise to persistent neural activity. We then discuss recent experiments that are beginning to test predictions from different potential mechanisms during WM. Finally, we discuss the types of future experimental and theoretical studies that appear most likely to resolve the long-standing questions posed above.

CELL-AUTONOMOUS BIOPHYSICAL MECHANISMS FOR GENERATING PERSISTENT ACTIVITY AND SIMPLE MEMORY-ENCODING CIRCUITS

Many central nervous system (CNS) neurons contain biophysical mechanisms that enable them to continue to discharge after the triggering input has subsided. One of the first critical insights to emerge from the first decade of intracellular recordings from acute brain slices was the discovery that individual neurons could generate all-or-none bursts of action potentials that prolonged responses to brief synaptic input, shown in hippocampal pyramidal cells and subsequently in other cell types (summarized in Traub & Jefferys 1994). These stereotyped discharges, comprising 3–6 spikes, closely resembled interictal discharges observed in animal models of epilepsy (Traub & Jefferys 1994, Wong et al. 1986).

The surprising insight that emerged from this early work was that a defining feature of at least one type of epileptic discharge was normally present in a subset of neurons that generated all-or-none bursts in nonepileptic animals. This initial experimental work led to a remarkably accurate set of predictions from realistic computer models of hippocampal networks (Miles & Wong 1987, Traub & Wong 1982, Wong et al. 1986) that a change in inhibitory circuits enabled interictal epileptic activity by allowing all-or-none bursts in one cell to percolate through the network. This classic work on burst transduction helped define a general theme in which regenerative biophysical currents are combined with recurrent synaptic connections to generate a collective network behavior.

As discussed below, this theme provides an important model for assimilating recent progress in understanding how another emergent network behavior—persistent firing during WM tasks—occurs. In this section, we review the three major classes of biophysical mechanisms that can enable persistent firing (voltage-gated Ca^{2+} and Na^{+} currents, inward currents that track intra-cellular Ca^{2+} , and Ca^{2+} -triggered long-term changes in neuronal excitability). As with interictal bursting, the origin of persistent firing in the cortical regions where it is most often studied (e.g. prefrontal and temporal neocortex, the hippocampus and entorhinal cortex) is likely to include both synaptic circuit and intrinsic biophysical components. The potential involvement of intrinsic conductances in generating persistent firing comes, in part, from the common observation that simple intracellular depolarization can initiate firing that outlasts the stimulus even when synaptic transmission is blocked pharmacologically (Fransén et al. 2006, Jochems & Yoshida 2013, Knauer et al. 2013, Navaroli et al. 2012, Pressler & Strowbridge 2006). Most of these studies involved adding exogenous modulators (typically drugs that activate muscarinic receptors, such as carbachol) that function to enhance excitability, perhaps mimicking the normal enhanced release of acetylcholine and other neuromodulators during heightened attention conditions associated with WM (see below). Neurons in subcortical areas, including regions involved in the control of eye movement (Arnold & Robinson 1997, Joshua & Lisberger 2015), also exhibit persistent firing rates in response to transient stimuli. The brainstem neurons involved in the vestibuloocular reflex appear to modulate their firing rates through network interactions, as direct depolarization of individual neurons (through intracellular recording pipettes) fails to evoke persistent discharges (Aksay et al. 2001).

The situation in cortical areas is less clear, with evidence (discussed below) supporting both the capability of individual neurons to fire persistently through biophysical mechanisms and the existence of recurrent excitatory synaptic connections among principal cells (Douglas & Martin 2004, Miles & Wong 1986)—proving the potential components for intrinsic- and synaptic circuit-mediated persistent activity. As the experimental and mathematical tools necessary to disambiguate these mechanisms are rapidly progressing, our primary aim is to provide a view of the evidence supporting different potential mechanisms as well as the new approaches that are beginning to yield insights into how and when they likely participate in memory-related sustaining firing modes.

Bistability Mediated by Voltage-Gated Inward Currents

One of the best-studied biophysical mechanisms for extending firing beyond the initiating stimulus is through plateau potentials mediated by regenerative Ca^{2+} or Na^+ currents. Plateau potentials in some types of CNS neurons appear to be generated by L-type voltage-gated Ca^{2+} channels (VGCCs) and can persist for many seconds (Otsuka et al. 2001, Perrier et al. 2002, Russo & Hounsgaard 1996). Even weak depolarizing inputs, such as a few summing excitatory postsynaptic potentials (EPSPs), can activate L-type VGCCs, leading to an inward Ca^{2+} current that depolarizes the neuron. This initial small Ca^{2+} current causes more VGCC activation, further depolarization, and eventually repetitive spiking. This biophysical mechanism takes advantage of the same regenerative properties that generate action potentials where inward currents amplify weak excitatory inputs. While the constellation of currents that generate action potentials includes outward currents, such as the delayed rectifier K^+ currents in neurons, that curtail the regenerative depolarizing response, the functional influence of these K^+ currents appears to be less in plateau-generating neurons. One potential mechanism for enabling regenerative plateau potentials, therefore, is to segregate clusters of inward and outward currents in different dendritic compartments, reducing the ability of outward currents also activated by synaptic inputs to abort regenerative plateau responses. Dendritic segregation of different subclasses of ionic currents also appears to be critical for the generation of all-or-none intrinsic burst responses (Goldman et al. 2003, Pinsky & Rinzel 1994).

The long duration of plateau potentials in spinal motoneurons likely reflects the minimal inactivation of L-type VGCCs during sustained depolarizations (Perrier et al. 2002). However, long-duration plateau potentials can also be mediated by inward currents that show pronounced inactivation, such as T-type Ca^{2+} channels and even voltage-gated Na^+ channels. In both cases, inactivation is incomplete, permitting a range of membrane potentials that can create steady-state inward currents through the generation of a window current. In thalamic neurons, the voltage range for triggering depolarization mediated by window currents through T-type Ca^{2+} channels is near the resting membrane potential (Crunelli et al. 2005, Hughes et al. 1999), enabling relatively weak synaptic stimuli to trigger persistent firing.

Recent work from Yamada-Hanff & Bean (2013) extended this model to explain long-lasting firing patterns observed in hippocampal neurons tested when intrinsic excitability is enhanced through pharmacological activation of cholinergic receptors. Whereas voltage-

gated Na⁺ channels show rapid and nearly complete inactivation following only a few milliseconds of strong depolarizing stimuli, the remaining persistent Na⁺ current can trigger prolonged firing. Like the run-away depolarization that underlies the action potential, the recruitment of subthreshold Na⁺ current is regenerative—the effects of a weak triggering depolarization are amplified by triggering intrinsic Na⁺ current, which then recruits more subthreshold Na⁺ current. This cycle continues until firing begins. Individual neurons can exhibit different firing modes likely supported, at least in part, by subthreshold Na⁺ current and reflecting different initial intrinsic physiological conditions (Kass & Mintz 2006).

One defining feature of persistent firing mediated by noninactivating inward currents is high sensitivity to hyperpolarizing input. In all three well-studied examples of this phenomenon (mediated by L-type channels, window currents through T-type Ca²⁺ channels, and subthreshold Na⁺ currents), interrupting persistent firing by injecting hyperpolarizing current abolishes firing rapidly (Otsuka et al. 2001). Once abolished, persistent firing does not resume if the initial triggering stimulus has subsided. This sensitivity to hyperpolarizing input [or to synaptic inhibition (Anderson & Strowbridge 2014)] distinguishes intrinsic persistent responses mediated by noninactivating voltage-gated currents from the other two major biophysical mechanisms supporting persistent firing: Ca²⁺-activated inward currents and modulation of slow intrinsic currents.

Calcium-Dependent Intrinsic Mechanisms

Persistent firing responses in many types of cortical neurons (including pyramidal cells in entorhinal cortex, hippocampus, and neocortex) can be triggered by depolarizing stimuli but likely involve a second messenger that activates the underlying depolarizing response. The primary evidence for an indirect pathway is the sensitivity of persistent firing to chelation of Ca²⁺ ions (Lei et al. 2014, Rahman & Berger 2011). Simple biophysical mechanisms, such as triggering regenerative inward Ca²⁺ or Na⁺ currents, are triggered directly by depolarizing stimuli and therefore should not be abolished by blockade of the intracellular signaling function of Ca²⁺ ions. In many neurons that exhibit intrinsic persistent activity [e.g., Blanes cells in the olfactory bulb (Pressler & Strowbridge 2006)], uncaging Ca²⁺ can depolarize the membrane potential, potentially triggering the same second-messenger responses as depolarizing stimuli that cause Ca²⁺ accumulations.

The identity of the Ca²⁺-dependent current that enables cortical neurons to remain depolarized following stimuli has remained elusive despite several decades of effort from many laboratories. One potential mechanism is a form of Ca²⁺-activated nonselective cation current (I_{CAN}) (Pace et al. 2007, Rubin et al. 2009). Although the molecular identity of the ion channel mediating this response is unknown, it appears to be gated by Ca²⁺ (or Ca²⁺ ions bound to an accessory protein), is permeable to Na⁺ and Ca²⁺, and is sensitive to flufenamic acid (FFA; Haj-Dahmane & Andrade 1999, Lei et al. 2014, Rahman & Berger 2011, Tahvildari et al. 2008, Zhang et al. 2011), a nonsteroidal anti-inflammatory agent. Several groups have hypothesized that a subtype of transient receptor potential (TRP) channels may mediate I_{CAN} responses, based on pharmacological studies (Lei et al. 2014, Reboreda et al. 2011, Petersson & Fransén 2012) and tests of peptides that interfere with TRP channel function (Yan et al. 2009, Zhang et al. 2011). However, thus far no direct evidence has

demonstrated that persistent firing and the associated afterdepolarization response are abolished in TRP channel knockout animals [although Lei et al. (2014) found partial blockade in TRPM5 knockout mice]. Given the wide diversity of TRP channels (Reboreda et al. 2011), fully testing this hypothesis may require generating mice that are null for many different TRP subunits.

There are two main challenges in determining the underlying mechanism of intrinsic persistent firing in cortical neurons. First, the most commonly used blocker, FFA, is nonspecific, with effects reported on a wide variety of ion channels, including K^+ channels (Guinamard et al. 2013). Blockade of persistent firing typically requires high concentrations of FFA, compounding concerns about its specificity. Second, the activation of a Ca^{2+} -dependent inward current such as I_{CAN} should be accompanied by an obvious decrease in input resistance. Although it is straightforward to monitor input resistance by injecting weak current test pulses or applying ramp protocols, interpreting the results from these experiments is difficult because the normal input resistance of most cortical neurons changes depending on the membrane potential (e.g., Yamada-Hanff & Bean 2013). In neocortical neurons tested under conditions in which depolarizing stimuli trigger persistent activity (cholinergic receptor activation), brief depolarizing stimuli triggered an apparent increase in input resistance—the opposite response expected from an I_{CAN} mechanism (Haj-Dahmane & Andrade 1996). However, the authors interpreted the resistance change as more likely consistent with an inwardly rectifying I_{CAN} mechanism, based on a failure of the response to reverse near the K^+ equilibrium potential and a sensitivity to manipulations of extracellular ions other than K^+ . More recent work has repeated the input resistance test and found an increase in input resistance even when compensating for expected change in input resistance during the afterdepolarization response (Cui & Strowbridge 2016, 2017), leaving open the possibility that alternative Ca^{2+} -triggered ionic mechanisms may be responsible for the afterdepolarization and persistent firing in neocortical neurons.

One intriguing aspect of the I_{CAN} hypothesis is that it explained the ability of entorhinal neurons to generate graded persistent activity, in which multiple depolarizing stimuli presented in succession each increased the frequency of persistent firing (Egorov et al. 2002, Fransen et al. 2006). I_{CAN} -mediated responses presumably continuously reflect changes in intracellular Ca^{2+} , as the underlying conductance is gated by Ca^{2+} . Although intracellular Ca^{2+} dynamics are complex, the I_{CAN} hypothesis predicts that second (and later) depolarizing inputs should increase the frequency of persistent firing as long as the stimulus is strong enough to trigger intracellular Ca^{2+} accumulation and I_{CAN} current is not saturated or inactivated.

There are two important caveats to the I_{CAN} hypothesis, beyond the absence of an established molecular identity for the ion channel mediating I_{CAN} . First, the classic study (Egorov et al. 2002) that demonstrated graded persistent firing used sharp microelectrodes instead of the now-ubiquitous whole-cell patch clamp recording method. Subsequent studies using patch clamp recordings did not observe graded persistent firing when the amplitude of the depolarizing step amplitude (Rahman & Berger 2011 in neocortex and Pressler & Strowbridge 2006 in the olfactory bulb) or step duration (Knauer et al. 2013 and Jochems & Yoshida 2013 in the hippocampus) was varied [although Navaroli et al. (2012) is an

exception], perhaps reflecting the role of a key intracellular signaling molecule that is lost through diffusion in whole-cell recordings. Sharp microelectrode recordings also introduce potential complications, primarily related to lower input resistance due to current leakage around the microelectrode, leaving open the question of which intrinsic firing mode (bistable or graded/multistable) is most likely to occur endogenously.

The second caveat is that when tested experimentally in both entorhinal and neocortical neurons, the frequency of persistent firing does not appear to track intracellular Ca^{2+} concentration. Pausing persistent firing long enough for intracellular Ca^{2+} concentrations to subside to near basal levels [for more than 4–5 s, given the approximately 1-s decay time constant for intracellular Ca^{2+} transient following muscarinic receptor activation (Cui & Strowbridge 2016, Rahman & Berger 2011)] should abolish firing because the underlying I_{CAN} -mediated depolarization should be strongly attenuated. When hyperpolarizing steps are applied after triggering persistent firing, they typically fail to abolish persistent firing, which resumes at only slightly lower frequencies after the hyperpolarization step ends [e.g., figure 1 in Fransén et al. (2006) with approximately 20-s firing pauses applied to entorhinal neurons]. Injecting hyperpolarizing pulses to pause firing for 5–10 s also fails to abolish persistent firing in neocortical neurons (Cui & Strowbridge 2016, Navaroli et al. 2012). But multiple studies also find that injecting very strong hyperpolarizing steps—more than required to simply pause firing—often does abolish, or strongly reduce, persistent firing (Egorov et al. 2002, Fransén et al. 2006, Knauer et al. 2013, Navaroli et al. 2012). Although the origin of the blocking effect with very strong hyperpolarization remains to be determined, failure of long-duration firing pauses to abolish persistent firing suggests that intracellular Ca^{2+} transients trigger a long-lasting change in intrinsic excitability. Rather than employing an I_{CAN} depolarizing response that continuously tracks intracellular Ca^{2+} concentration, the biophysical mechanism underlying persistent activity may reflect a covalent modification (e.g., phosphorylation of an ion channel) that effectively decouples the dynamics of the triggering stimulus from the resultant persistent firing.

If intrinsic persistent firing reflects a long-lasting change in excitability, how does it occur? Hasselmo and colleagues (Fransen et al. 2006) hypothesized that the channels mediating I_{CAN} responses could exist in two states that differ in conductance. Similar to the auto-activation mechanism proposed to explain state changes in calmodulin-dependent protein kinase II activity (Lisman & Goldring 1988), this model proposes that the Ca^{2+} influx from the triggering stimulus not only opens I_{CAN} channels but also initiates a biochemical cascade that makes the channels generate larger currents in response to Ca^{2+} , facilitating long-lasting changes in excitability (Fransen et al. 2006). Because the channels mediating responses attributed to I_{CAN} have not been determined, this proposal has remained untested, although there are phosphorylation sites on candidate TRP channels that affect their function (Zhang et al. 2011) and could mediate long-lasting changes in excitability. What is still missing, though, are molecular or more specific pharmacological experiments that link TRP channels definitively to I_{CAN} and persistent firing.

Although less often studied in relation to persistent firing, a reduction in a standing (or leak) K^+ current also could mediate a long-lasting change in excitability and would explain the increase in apparent input resistance observed following depolarizing stimuli in cortical

neurons (Cui & Strowbridge 2016, 2017; Haj-Dahmane & Andrade 1996). Persistent firing is often studied in experimental conditions in which some leak K^+ current has been reduced, as muscarinic receptor agonists attenuate I-M (Jentsch 2000). Transient Ca^{2+} accumulations could trigger a further decrease in I-M or another type of K^+ current active near the resting potential, causing depolarization and persistent firing. One potential candidate leak K^+ channel that has been shown to be attenuated through a Ca^{2+} -dependent signaling cascade is ether-à-go-go-related gene channel (ERG) (Cockerill et al. 2007, Pessia et al. 2008). Recent work demonstrated that three different ERG blockers abolish both the stimulus-triggered increase in apparent input resistance and persistent firing in neocortical pyramidal cells (Cui & Strowbridge 2016, 2017). The slow activation of ERG currents may contribute to the normal slowing of firing in response to depolarizing steps (Pessia et al. 2008). Strong depolarizing stimuli combined with muscarinic receptor activation appears to trigger a long-lasting attenuation of ERG-mediated K^+ current, leading to less spike frequency adaptation during the stimulus and a subsequent 10–15-s period of hyperexcitability. Relatively little is known about ERG function in CNS neurons in vivo, which is difficult to probe by applying ERG blockers systemically because of the critical role ERG currents play in cardiac cells. Determining the functional role of ERG currents will likely entail using either molecular approaches or focal delivery of specific blockers. Unlike the relatively nonspecific blocker FFA, which also attenuates ERG currents (Guinamard et al. 2013), highly specific blockers are available for ERG channels, including a potent peptide toxin made endogenously by scorpions (Nastainczyk et al. 2002).

Although significant gaps remain in our understanding of the biophysical details, recent work has identified three major classes of mechanisms that could support intrinsic persistent activity (voltage-gated Na^+ and Ca^{2+} currents, Ca^{2+} -tracking nonselective cation currents, and long-lasting state changes in either K^+ or cation currents that function to depolarize neurons). We currently have much better experimental methods for determining which of these types of mechanisms operate in specific neurons than we have for determining the actual underlying ion channel. Persistent activity resulting from voltage-gated ion channels is generally easily abolished by even brief hyperpolarizations and is not strongly affected when intracellular Ca^{2+} transients are attenuated using Ca^{2+} chelators. Repetitive depolarizing stimuli and long-duration hyperpolarizing steps can be used to differentiate between mechanisms that rely on intracellular Ca^{2+} signaling to modulate ion channels rapidly (e.g., direct modulation of I_{CAN} currents by Ca_2^+) or to effect a state change in excitability (e.g., via covalent modulation of I_{CAN} or leak K^+ channels).

Memory Functions Created by Embedding Neurons with Intrinsic Bistability in Feedforward Networks

When viewed from the single-cell perspective, neurons that generate cell-autonomous persistent firing represent simplistic, binary memory. However, once embedded in even simple feedforward networks, the resulting system can generate circuit-specific representations of biological information reliably. The memory-encoding function of these circuits relies on two critical factors: divergent projections from the intrinsically bistable neurons onto a population of downstream neurons, and nonuniform activation of bistable neurons by orthodromic synaptic input patterns (Figure 1). These features enable different

inputs presented to the network to trigger distinctive changes in the spontaneous synaptic tone. With a large enough fan-out from the bistable neurons, the memory of the input pattern presented can be deduced by monitoring the rate of spontaneous EPSPs in a small subset of downstream neurons—often a much easier experimental task than locating and monitoring the actual memory-originating bistable neurons. If the bistable neurons are excitatory, the stimulus-specific changes in synaptic tone could also generate distinctive patterns of action potentials that would be revealed using standard extracellular or optical recording methods.

The first biological example of this type of memory encoding was demonstrated in the dentate gyrus, an evolutionary old brain region containing a small population of glutamate neurons that generate VGCC-dependent plateau potentials [semilunar granule cells (SGCs)] (Williams et al. 2007). SGCs are excited by perforant path inputs originating from entorhinal cortex and project to mossy cells and interneurons contained in an adjacent subfield (the dentate hilus; Larimer & Strowbridge 2010, Williams et al. 2007), enabling different stages of this feedforward circuit to be assayed independently. Stimulating different subgroups of perforant pathway axons triggered plateau potentials in different subgroups of SGCs, leading to EPSP barrages in downstream hilar neurons that could be assayed up to 10 s later to determine which pathway was activated (Hyde & Strowbridge 2012, Larimer & Strowbridge 2010). Because excitatory hilar mossy cells only very rarely form recurrent connections (Larimer & Strowbridge 2008), this form of short-term memory almost certainly does not involve any of the recurrent circuit topologies described below. Despite its simplicity, this intrinsic/feedforward network hybrid can encode contextual information, such as the order of temporal sequences (Hyde & Strowbridge 2012). In more complex cortical brain regions such as neocortex, the physiology of neurons with the capacity for intrinsic bistability is under modulatory control, and those cells are intermixed with nonbistable cells. Excitatory neo-cortical pyramidal cells also form relatively abundant recurrent connections (e.g., Morishima et al. 2011), making it difficult to know if similar stimulus-specific changes in synaptic tone result from dentate-like intrinsic/feedforward hybrid networks or from more complex circuit architectures.

NEURAL CIRCUIT MECHANISMS FOR GENERATING PERSISTENT REPRESENTATIONS

Neural circuit mechanisms capable of generating persistent representations have been reviewed extensively elsewhere (Brody et al. 2003, Chaudhuri & Fiete 2016, Compte 2006, Durstewitz et al. 2000, Hasselmo & Stern 2006, Major & Tank 2004, Wang 2001), and so we provide a brief summary here. The key challenge for these models is that synaptic responses in most cortical neurons (i.e., those not expressing the persistent currents described above) decay within 10–20 ms, and so these models seek patterns of connectivity within the circuit that will extend these very short-term representations such that they last for hundreds of milliseconds, up to tens of seconds, or more. There are two main ways to achieve this. First, the circuits can have attractor states, leading to activity patterns that are stable in time. Alternatively, the circuits can display time-varying neural activities for which certain combinations of neural firing rates nevertheless remain fairly constant and thus can represent the inputs persistently. For example, if some neurons' firing rates increase over

time while others decrease, the summed firing rates of the population can remain stable even as the individual neural firing rates vary. Both of these classes of mechanisms have some experimental support and are thus apparently biologically plausible. In what follows, we briefly describe each mechanism and then discuss the experimental evidence for, or against, said mechanism.

Networks with Attractor States—Theoretical Ideas

The classic way to achieve persistent representations is for the neural population to have specific patterns of activity that are reinforced by the synaptic connectivity in the network (Hopfield 1982); we refer to these as privileged activity patterns. These activity patterns, once generated, are stably reinforced by the reverberant synaptic inputs to the active neurons, and the network structure is such that, from any initial pattern of activation, the population activities evolve gradually toward one of these privileged patterns (Figure 2). We first consider models with discrete numbers of privileged activity patterns and discuss models with continua of such patterns below.

As a simple example, consider a network where neurons are divided into two groups (A and B). Neurons within each group strongly excite each other and strongly inhibit neurons within the other group. If the network initially has more active neurons in group A than in group B, mutual excitation between group A neurons drives up activity in that group, and inhibition drives down activity in group B neurons. Similarly, if the network initially has more activity in group B than A, the network evolves so that most or all group B neurons are active, and few group A neurons are. Accordingly, this network has two privileged activity patterns (“A on” or “B on”), and the one that gets activated depends on the initial state of the network. This initial state can be determined by external inputs. The privileged patterns are known as attractor states because the network activity naturally evolves towards them. Although our illustrative example has two attractor states, with appropriately chosen patterns of connectivity, neural circuits can have as many attractor states as they do neurons (Hillar et al. 2012). Attractor networks of this sort have, from the outset, been studied with spiking (or binary) neuron models (see the sidebar titled Different Classes of Computational Neuron Models for descriptions of the different types of neuron models) (Amit & Brunel 1997, Hopfield 1982).

DIFFERENT CLASSES OF COMPUTATIONAL NEURON MODELS

Spiking neuron models are the most biologically realistic of the computational neuron models considered here. They describe the evolution of the neuron’s membrane potential over time and the corresponding generation of action potentials. In the simplest such models—known as integrate and fire models—a single membrane potential variable (presumably the somatic voltage) is considered, and whenever that value crosses the spiking threshold, the neuron emits an action potential. Thereafter, the membrane potential is reset to its resting value, from which it can subsequently be depolarized by appropriate synaptic inputs. This neuron model is used, for example, in the bump attractor network model of Compte et al. (2000). More complex (and correspondingly more realistic) spiking neuron models compute the membrane potential at many different

parts of the cell and include a range of different ionic currents and cellular compartments. An example of this type of model can be found in Traub & Wong (1982).

Instead of explicitly modeling the dynamical evolution of the neurons' states, binary neuron models take—at each time point in the simulation—a weighted combination of their inputs and compare the results to a threshold. If the input is superthreshold, the neuron emits a spike (giving an output value of 1), and if the input is subthreshold, the neuron does not spike (giving an output value of 0). These models capture the discrete nature of interneuronal communication but do not capture any of the temporal properties of the way neurons integrate their inputs because all inputs are considered to arrive at the cell at the same time. This type of neuron model is used in the discrete attractor network of Hopfield (1982).

In rate-based models, instead of modeling the sequences of action potentials emitted by the neurons, the models instead describe each neuron by a continuous-valued rate variable that changes over time in response to inputs from other neurons. These continuous variables are meant to describe the firing rates of the modeled neurons. Rate-based models do not capture the discrete nature of the interneuronal communication via action potentials and typically do not describe any of the complex biophysical properties of the neurons. Nonetheless, they are more amenable to mathematical analysis than are the more complex neuron models described above, and so they are often used in theoretical studies. Examples include Seung (1996) and Druckmann & Chklovskii (2012).

The above discussion describes networks with discrete numbers of attractor states, which means that they can store only finite numbers of different inputs. In practice, especially for spatial WM tasks, it is desirable to store continuous-valued information (the coordinates of the stimulus): Other WM tasks also involve remembering continuous-valued information (Romo et al. 1999, Zaksas & Pasternak 2006). This motivated the development of models with continuous attractors that could stably maintain a range of input values (Burak & Fiete 2009, Compte et al. 2000, Ganguli et al. 2008, Seung 1996). Here, instead of a discrete set of privileged attractor states, the network connectivity is such that there is a continuum of attractor states. As an illustrative example, we consider the bump attractor models (Amari 1977, Compte et al. 2000). The idea here is that the spatial position that the animal needs to remember is encoded by activating the neurons that respond preferentially to that position. When the neurons are sorted by their preferred positions (so that, for visualization, neurons with similar preferred positions are placed near each other), this activation corresponds to a bump of activity: The bump position describes the stimulus location that is being remembered. Connectivity in the network is such that neurons with similar preferred positions excite each other, and those with dissimilar preferred positions inhibit each other. This pattern of connectivity sustains the bump of activity after it is triggered. Because the activity bump could be at any different position, this is an example of the continuous class of attractors. Whereas early continuous attractor models used rate-based neurons (Amari 1977, Seung 1996), more recent work has shown that the same ideas apply to the more realistic spiking neuron models (Compte et al. 2000).

In practice, much fine-tuning of the network connectivity is needed in order to form continuous attractors. For example, the bump attractors described above require that the connectivity between cells be translation invariant, such that all pairs of cells with the same spacing between their preferred positions have the same strength of synaptic connectivity. Without this symmetry, the bump is not stably maintained at all positions and instead tends to drift toward the positions corresponding to the most strongly excited (and least strongly inhibited) cells. Other types of continuous attractors, such as the line attractor model of Seung (1996), require that the matrix describing connectivity between pairs of cells has a very precise mathematical property (namely, the largest eigenvalue must be equal to 1). When the largest eigenvalue is less than 1, the network activity gradually decays, instead of being stably maintained, and so the memory information is lost over time. And when the largest eigenvalue exceeds 1, the network activity continually increases, eventually running away to extremely large values.

As a solution to this fine-tuning problem, networks have been proposed with multiple, finely spaced discrete attractors (for which the required synaptic weight and/or wiring precision is less) and that combine intrinsic neuronal persistence with network mechanisms, so as to improve the system's robustness (Goldman et al. 2003, Kilpatrick et al. 2013, Koulakov et al. 2002). Networks that incorporate bistable neurons can tolerate changes in synaptic weights of up to 20% without losing memory function. Other work has explored plasticity mechanisms that adjust the network connectivity dynamically to maintain the required precision (MacNeil & Eliasmith 2011, Renart et al. 2003).

Another potential problem faces continuous attractor networks: perfectly tuned networks can be sensitive to random fluctuations. For example, synaptic failures (Branco & Staras 2009), ion channel fluctuations (Goldwyn & Shea-Brown 2011, White et al. 2000), and ongoing noisy inputs from other brain areas (Churchland et al. 2010) or from the sense organs (Field et al. 2005, Zylberberg et al. 2016a) can cause random changes in the inputs experienced by cortical neurons. For perfectly tuned bump attractors, this noise will cause the activity bump to move randomly, changing the remembered position; similar statements apply to other types of continuous attractors. Here again, networks with multiple finely spaced discrete attractors are more robust than true continuous attractors (Cain et al. 2013, Goldman et al. 2003, Kilpatrick et al. 2013, Koulakov et al. 2002) because relatively small perturbations that do not move the activity between adjacent attractor states will not have a long-lasting effect.

Finally, it is important to note that because attractor networks rely on reverberant interactions within the network, they are sensitive to the different timescales of excitation relative to inhibition (Wang 1999): If reverberant excitation is much faster than reverberant inhibition, the neural firing rates could become very large before inhibition has a chance to catch up. Theoretical work shows that, owing to the relatively slow dynamics of NMDA receptor synapses, NMDA receptor-mediated excitation could be critical to the generation of stable persistent activity in neural circuits (Tegnér et al. 2002, Wang 1999). Other mechanisms might also exist that could delay or slow the buildup of persistent activity through the dynamics of modulations associated with the high attention states that are often

behaviorally linked with WM. Understanding how neuromodulation-related dynamics interplay with circuit ramp-up is likely to be defined in the near future using optogenetics.

Networks with Attractor States—Experimental Evidence

Recent studies of the sensory systems, and of the hippocampus, have made observations that are consistent with the existence of a discrete set of attractor states (Bathellier et al. 2012, Niessing & Friedrich 2010, Wills et al. 2005). Over large neural populations, activity patterns tended to cluster into a small number of states (Bathellier et al. 2012), and during continuous (smooth) changes in the environment, neural activities tended to switch abruptly between discrete states (Niessing & Friedrich 2010, Wills et al. 2005) as opposed to showing a smooth variation.

Although the above studies were not done explicitly in the context of WM, other work used optogenetics to more directly test whether discrete attractors are implicated in WM. Kopec et al. (2015) used optogenetics to transiently inactivate neurons in the FOF of mice at different times during the delay period of a memory-guided orienting task: Mice could turn left (L) or right (R) after the delay period to report their memory. Inactivation during the early phases of the delay period had a larger effect than inactivation in the later portion of the delay period, leading the authors to argue in favor of a discrete attractor mechanism. Specifically, Kopec et al. (2015) postulate two attractor states, one corresponding to each possible choice. The intuition is that, at the start of the delay period, neural activities are near the boundary between the states that will evolve either to the L activity state or the R one. Consequently, small perturbations at this early stage can push the neural activities across the boundary, changing the subsequent behavioral outcome. For contrast, later in the delay period, after the neural activities have evolved away from that boundary, toward the L or R states, the same magnitude of perturbation is unable to push the population activity state across the boundary. Although this experiment does show clearly that something about the upcoming decision crystallizes within the network during the delay period, and the data are recapitulated by a computational model with two attractor states, the data nevertheless do not prove unambiguously that the neural circuit has an attractor structure. For example, the circuit could contain bistable neurons of the types discussed above, for which the persistent firing states take a while to be activated. In that case, perturbation early in the delay period would affect the cells before the persistent firing states get fully activated and would thereby have a much larger effect than would later perturbation, after the persistent firing has been established.

With regards to continuous attractor models, observations of grid cells in the entorhinal cortex show what appear to be continuous attractor dynamics (Yoon et al. 2013). More closely related to WM, observations have been made in monkey lateral intraparietal cortex during a decision-making task that are consistent with the neural activities converging towards a (continuous) line attractor (Ganguli et al. 2008).

Moreover, Wimmer et al. (2014) argued that data from monkey PFC during a spatial WM task support the bump attractor hypothesis described above. Notably, even over sets of trials in which the same stimulus was presented to the animal, the trial-by-trial differences in the delay period neural activity were predictive of the differences in the animals' subsequent

behavioral outcomes. Furthermore, the levels of trial-to-trial variability in the delay period activities, and the correlations between different neurons' delay period activities, depended systematically on the location of the presented stimulus. All these effects could be predicted by the bump attractor model, wherein the animal's memory is encoded by a bump of neural activity that moves about during the delay period. This drift in the activity bump is presumably caused by either task-irrelevant network inputs that differ slightly between trials or random noise (Churchland et al. 2010, Faisal et al. 2008, White et al. 2000). Because the recurrent bump attractor model can successfully predict both the relation between trial-specific neural activity fluctuations and behavioral deviations, and the statistical properties of the delay period activities, the Wimmer et al. (2014) results are perhaps the strongest argument for persistent activity arising from precisely wired recurrent networks as opposed to simpler feedforward networks that include neurons with intrinsic persistence.

Models with Persistent Representations Despite Time-Varying Neural Firing Patterns— Theoretical Ideas

The attractor models described above all have one thing in common: In response to an input, they settle into their attractor activity patterns and stay there. In other words, these models predict that, after an initial transient, neural activities should be constant during the delay period of a WM task. Although some neurons have been observed with nearly constant activities during these delay periods (Constantinidis et al. 2001, Funahashi et al. 1989, Fuster & Jervey 1981, Myashita & Chang 1988), other observations show more variable delay period activity (Barak et al. 2010, Brody et al. 2003, Shafi et al. 2007, Zaksas & Pasternak 2006). Moreover, observations of persistent activity in feedforward neural circuits (Hyde & Strowbridge 2012, Larimer & Strowbridge 2010) bring up the question of whether recurrent circuitry—a key feature of the attractor models—is necessary in the neural circuits that generate persistent activity. These experimental findings motivate new circuit models that relax the requirements of the attractor models while maintaining the ability to generate persistent representations.

The first such class of models removes the need for constant delay period neural activity by exploiting the fact that, even if individual neurons' firing rates fluctuate, appropriate combinations of those firing rates can remain constant. Recent work by Druckmann & Chklovskii (2012) shows that recurrent circuits can indeed have this property. This model takes advantage of the likely distributed nature of memory representations in cortical networks. Whereas some cells are strongly tuned to only one (remembered) stimulus feature, often, individual cells are selective for many features, and multiple cells code for each feature (Rigotti et al. 2013). The Druckman & Chklovskii (2012) model exploits this redundant coding to create stable representations by combining the outputs of different single units into weighted sums (Figure 2). For these combinations of neural firing rates to remain constant, this approach requires fine-tuned network connectivity and will not work with random connectivity. [Specifically, the matrix describing recurrent connectivity in the network must have an eigenvalue very near to 1: The corresponding eigenvector describes the weighting of the neurons in the weighted sum that remains constant even as the neural activities vary. This is the same tuning requirement displayed by the original line attractor

networks (Seung 1996).] Similar models have been studied elsewhere in the context of WM and neural integrators (Boerlin et al. 2013, Murray et al. 2016, Schwemmer et al. 2015).

Other studies asked whether recurrent connectivity is needed at all. That work revealed that feedforward chains can extend the representation beyond the timescale of single-neuron activation (Ganguli & Latham 2009, Goldman 2009). The idea is that after the stimulus goes away, excitation continues to propagate through the chain from one cell (or cell group) to the next, in a manner reminiscent of the older concept of synfire chains (Abeles 1991). So long as this chain is long enough, the input can be maintained for an extended period (intuitively, for a period of approximately $N \tau$, where N is the length of the chain, and τ is the single-neuron time constant). These feed-forward chains can be embedded within recurrently coupled networks (Goldman 2009, Trengove et al. 2013): Even in circuits with recurrent connectivity, there can be “hidden” feedforward chain structures. These can be revealed by appropriate mathematical analysis of the connectivity matrix (Goldman 2009) or by observing the sequentially activated neurons. Neuronal networks can learn via synaptic plasticity to generate the sequential activations (Rajan et al. 2016, Savin & Triesch 2014).

Finally, whereas all the above-reviewed mechanisms require specially constructed network connectivity, another body of work has investigated the functions that can be performed by networks of randomly (or nearly randomly) connected neurons. That work shows that, so long as there is an appropriate balance between excitation and inhibition in the circuit, information about the network’s inputs will be maintained after those inputs go away (Bertschinger & Natschläger 2004, Buonomano & Maass 2009, Laje & Buonomano 2013, Maass et al. 2002, Singh & Eliasmith 2006, Toyozumi & Abbott 2011). Similar to the work of Druckmann & Chklovskii (2012), an appropriate readout of the neural activities can thus yield a stable representation, even if the neurons’ activities vary in time (Barak et al. 2013).

Models with Persistent Representations Despite Time-Varying Neural Firing Patterns— Experimental Evidence

Experiments in brain slices and in vivo provide supporting evidence for models of dynamical neural activities, in which certain combinations of firing rates nevertheless remain constant (Li et al. 2016, Murray et al. 2016, Zylberberg et al. 2016b). Of particular note is the recent work by Murray et al. (2016), who compared neural recordings from monkey PFC during the delay period of both a spatial WM task and a parametric WM task to the predictions of several different computational models. Of the models considered, the constant combination of firing rates model (Druckmann & Chklovskii 2012) was found to be most consistent with the electrophysiology observations. Relatedly, Li et al. (2016) analyzed delay period activity to identify weighted sums of neural activity that best predict the upcoming behavioral report (similar to the notion of the coding line in Figure 2c). They then used optogenetics to inhibit neural activity for part of the delay period of their WM task. Surprisingly, after release of optogenetic inhibition, the behaviorally relevant weighted sums of neural activity recovered to resemble their values on trials without optogenetic perturbation; other combinations of neural activities (those that did not appear to encode task-relevant information) showed much less recovery.

At the same time, other work in rodents shows instead sequentially activated neurons during the delay period (Harvey et al. 2012, MacDonald et al. 2011, Pastalkova et al. 2008), providing support for the synfire chain-like model of Goldman (2009). However, other mechanisms (besides the feedforward chains) can also give rise to sequentially activated neurons, and so the experimental findings are not unambiguous. For example, consider a set of bistable neurons, as reviewed above, all with different durations of their triggered periods of persisting spiking. Downstream cells that respond to the turning off of these cells' persistent firing will show sequential activation, without requiring chain-like connectivity.

EXPERIMENTS SUGGESTING PRINCIPLES UNDERLYING THE ROLE OF PERSISTENT ACTIVITY IN WORKING MEMORY

The experiments above hint at specific mechanisms underlying persistent activity in neural circuits. Two other studies deserve discussion here precisely because rather than pointing at specific mechanisms, they may reveal more general principles underlying WM.

First, Li et al. (2016) measured both behavioral responses and neural activities after optogenetic perturbation during the delay period of a WM task. In this task, mice reported their choice by licking either the left or right port at the end of the delay period. Neural recordings and optogenetic perturbations were done in ALM, an area known to be involved in planning licking movements. Transient optogenetic silencing in one hemisphere of ALM abolished activity in that hemisphere's ALM. After the perturbation, activity recovered to be similar to that observed on trials without optogenetic perturbation. This observation casts doubt on the notion that there is a simple recurrent network, localized within one ALM hemisphere, that sustains the persistent activity: In that case, the activity would be unlikely to rebound after the end of the optogenetic manipulation. Instead, the data suggest that there are either multiple redundant circuits, extending beyond the scale of a single ALM hemisphere, or intrinsic mechanisms—such as the Ca^{2+} -dependent ones reviewed above—that can restore the persistence after optogenetic manipulation has ended.

Supporting the notion of multiple redundant circuits, the activity rebound following optogenetic perturbation was not observed after bilateral ALM silencing (both hemispheres) or unilateral ALM silencing in animals with severed corpus callosum. That redundancy across functional modules may be key to robust WM function. At the same time, as demonstrated by Li et al. (2016), many different WM mechanisms within individual ALM hemispheres can account for the observations, so long as the two ALM hemispheres both redundantly encode the memory information and share it via callosal projections.

Second, Liu et al. (2014) investigated the delay period activity in mouse PFC, both while the mice were learning the WM task and after the animals were well trained. They found that, during learning, PFC neurons' delay period firing rates were more strongly modulated than they were after the animals were well trained. Those neurons whose firing rates increased during the delay period (relative to the cue period) had higher delay period activity during learning than after the animals were well trained, and those neurons whose firing rates decreased during the delay period (relative to the cue period) had lower delay period activity during learning, as compared to when the animals were well trained. These observations

suggest that PFC delay period activity might be more important in learning than in performing well-trained WM tasks. Further supporting that assertion, Liu et al. (2014) performed optogenetic manipulations in the PFC during the delay period. Similar to the neural recordings described above, they performed these perturbations either during the experimental sessions in which the mice were learning the task, or during those experimental sessions after the prolonged learning period, when the mice were well trained. They found that manipulations of PFC delay period activity had a profound effect on the animals' abilities to learn the task, but that after learning, the task performance was robust to the optogenetic perturbations. These surprising observations emphasize that not all persistent delay period activity is task relevant, and so caution is warranted in relating delay period activity to WM. Moreover, the Liu et al. (2014) findings imply different roles for delay period activity under different contextual conditions (novel versus familiar task). It will be interesting going forward to understand how learning alters the mechanisms that support WM.

FUNCTIONAL CONSTRAINTS ON PERSISTENCE-GENERATING MECHANISMS

The preceding sections of this review focused on our current understanding of the palette of potential biophysical and circuit mechanisms that could generate persistent activity. Going forward, an important task will be to explore how these mechanisms function in vivo, during normal behaviors such as WM tasks, and how they might be involved in clinically important pathologies. Understanding how persistently active networks function involves more than testing for candidate mechanisms but will likely need to include research on how modulatory systems enable persistence (e.g., how cholinergic receptor activation during periods of heightened attention affects intrinsic excitability) and how the dynamics of persistent spiking responses are controlled by inhibitory synaptic input and biochemical feedback pathways. Further challenges are to understand how tightly controlled persistent networks can remain stable under a wide variety of neuromodulatory conditions, how synaptic plasticity can carry out the formation and maintenance of WM networks, and how these networks remain robust against synaptic failures and other biological sources of variability.

The first challenge is that modulators can change the effective connectivity of a neural circuit substantially. For example, blocking I-M-type K^+ currents (by activating muscarinic receptors) can increase the input resistance of neurons en masse. That strengthens the effective synaptic connections between those cells, allowing the same input currents to cause larger changes to the membrane potentials. This means that a network that was previously fine-tuned to form a continuous attractor will no longer do so. Even for the nonattractor models wherein persistent representations are formed despite time-varying neural activities (Druckmann & Chklovskii 2012), fine-tuned connectivity is still required. Thus, it remains unclear how these circuits would maintain function in the broad spectrum of neuromodulatory environments likely present in vivo.

Relatedly, synaptic plasticity changes the strengths of interneuronal connections constantly (Abbott & Nelson 2000). Unless the changes caused by the synaptic plasticity support the

WM function (i.e., retune the network dynamically to better generate persistent activity), then that function may be lost quickly. Consequently, an important line of inquiry, with some advances already made (Litwin-Kumar & Doiron 2014, Renart et al. 2003, Savin & Triesch 2014, Zenke et al. 2015), is to understand how ongoing synaptic plasticity can create reliable and long-lasting WM networks.

Finally, although theoretical models typically contain perfectly reliable synapses, in practice, biology offers few synapses with this level of robustness. Instead, most synapses (especially those in cortex) are characterized by high failure rates (Branco & Staras 2009). Because the persistent activity networks demand relatively precise patterns of synaptic connection, and the biological variability means that effective connectivity changes rapidly and seemingly randomly, it remains somewhat mysterious how the WM networks can function in the real biological setting. Future work that addresses this issue will be most informative.

Acknowledgments

We thank Edward Cui, Rodrigo Andrade, and Albert Compte for helpful discussions related to this review. This work was supported by US National Institutes of Health grant R01-DC04285 to B.W.S. and by a Canadian Institute for Advanced Research Azrieli Global Scholar Award to J.Z.

Glossary

WM	working memory
PFC	prefrontal cortex
ALM	anterior lateral motor cortex
FOF	frontal orienting fields
SC	superior colliculus
CNS	central nervous system
VGCC	voltage-gated Ca ²⁺ channel
EPSP	excitatory postsynaptic potential
I_{CAN}	Ca ²⁺ -activated nonselective cation current
FFA	flufenamic acid, an anti-inflammatory drug that also blocks a wide variety of ion channels
TRP	transient receptor potential
I-M	M-current, a type of noninactivating K ⁺ current that is attenuated by signaling cascades initiated by muscarinic receptors
ERG	ether-à-go-go-related gene channel

LITERATURE CITED

- Abbott LF, Nelson SB. Synaptic plasticity: taming the beast. *Nat Neurosci.* 2000; 3:1178–83. [PubMed: 11127835]
- Abeles, M. *Corticonics: Neural Circuits of the Cerebral Cortex.* Cambridge, UK: Cambridge Univ. Press; 1991.
- Amari SI. Dynamics of pattern formation in lateral-inhibition type neural fields. *Biol Cybern.* 1977; 27:77–87. [PubMed: 911931]
- Amit DJ, Brunel N. Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cereb Cortex.* 1997; 7(3):237–52. [PubMed: 9143444]
- Anderson RW, Strowbridge BW. Regulation of persistent activity in hippocampal mossy cells by inhibitory synaptic potentials. *Learn Mem.* 2014; 21(5):263–71. [PubMed: 24737918]
- Arnold DB, Robinson DA. The oculomotor integrator: testing of a neural network model. *Exp Brain Res.* 1997; 113:57–74. [PubMed: 9028775]
- Aksay E, Gamkrelidze G, Seung HS, Baker R, Tank DW. In vivo intracellular recording and perturbation of persistent activity in a neural integrator. *Nat Neurosci.* 2001; 4:184–93. [PubMed: 11175880]
- Barak O, Sussillo D, Romo R, Tsodyks M, Abbott L. From fixed points to chaos: three models of delayed discrimination. *Prog Neurobiol.* 2013; 103:214–22. [PubMed: 23438479]
- Barak O, Tsodyks M. Working models of working memory. *Curr Opin Neurobiol.* 2014; 25:20–24. [PubMed: 24709596]
- Barak O, Tsodyks M, Romo R. Neuronal population coding of parametric working memory. *J Neurosci.* 2010; 30(28):9424–30. [PubMed: 20631171]
- Bathellier B, Ushakova L, Rumpel S. Discrete neocortical dynamics predict behavioral categorization of sounds. *Neuron.* 2012; 76(2):435–49. [PubMed: 23083744]
- Benna MK, Fusi S. Computational principles of synaptic memory consolidation. *Nat Neurosci.* 2016; 19:1697–706. [PubMed: 27694992]
- Bertschinger N, Natschläger T. Real-time computation at the edge of chaos in recurrent neural networks. *Neural Comput.* 2004; 16(7):1413–36. [PubMed: 15165396]
- Bisley JW, Zaksas D, Droll JA, Pasternak T. Activity of neurons in cortical area MT during a memory for motion task. *J Neurophysiol.* 2004; 91:286–300. [PubMed: 14523065]
- Boerlin M, Machens CK, Denéve S. Predictive coding of dynamical variables in balanced spiking networks. *PLOS Comput Biol.* 2013; 9(11):e1003258. [PubMed: 24244113]
- Branco T, Staras K. The probability of neurotransmitter release: variability and feedback control at single synapses. *Nat Rev Neurosci.* 2009; 10(5):373–83. [PubMed: 19377502]
- Brody CD, Hernández A, Zainos A, Romo R. Timing and neural encoding of somatosensory parametric working memory in macaque prefrontal cortex. *Cereb Cortex.* 2003; 13(11):1196–207. [PubMed: 14576211]
- Buonomano DV, Maass W. State-dependent computations: spatiotemporal processing in cortical networks. *Nat Rev Neurosci.* 2009; 10(2):113–25. [PubMed: 19145235]
- Burak Y, Fiete IR. Accurate path integration in continuous attractor network models of grid cells. *PLOS Comput Biol.* 2009; 5:e1000291. [PubMed: 19229307]
- Cain N, Barreiro AK, Shadlen M, Shea-Brown E. Neural integrators for decision making: a favorable tradeoff between robustness and sensitivity. *J Neurophysiol.* 2013; 109:2542–59. [PubMed: 23446688]
- Chaudhuri R, Fiete I. Computational principles of memory. *Nat Neurosci.* 2016; 19(3):394–403. [PubMed: 26906506]
- Churchland MM, Yu BM, Cunningham JP, Sugrue LP, Cohen MR, et al. Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat Neurosci.* 2010; 13:369–78. [PubMed: 20173745]
- Cockerill SL, Tobin AB, Torrecilla I, Willars GB, Standen NB, Mitcheson JS. Modulation of hERG potassium currents in HEK-293 cells by protein kinase C. Evidence for direct phosphorylation of pore forming subunits. *J Physiol.* 2007; 581(2):479–93. [PubMed: 17363390]

- Compte A. Computational and in vitro studies of persistent activity: edging towards cellular and synaptic mechanisms of working memory. *Neuroscience*. 2006; 139(1):135–51. [PubMed: 16337341]
- Compte A, Brunel N, Goldman-Rakic PS, Wang X-J. Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cereb Cortex*. 2000; 10(9):910–23. [PubMed: 10982751]
- Constantinidis C, Franowicz MN, Goldman-Rakic PS. The sensory nature of mnemonic representation in the primate prefrontal cortex. *Nat Neurosci*. 2001; 4(3):311–16. [PubMed: 11224549]
- Crunelli V, Tóth TI, Cope DW, Blethyn K, Hughes SW. The ‘window’ T-type calcium current in brain dynamics of different behavioural states. *J Physiol*. 2005; 562(1):121–29. [PubMed: 15498803]
- Cui, ED., Strowbridge, BW. Mechanisms regulating persistent spiking in rodent neocortical neurons in vitro. Presented at Soc. Neurosci. Annu. Meet., Poster 590.19; Nov. 12–16; San Diego. 2016.
- Cui, ED., Strowbridge, BW. Modulation of Ether-à-go-go related gene (ERG) current governs intrinsic persistent activity in rodent neocortical pyramidal cells. 2017. bioRxiv 156075. <https://doi.org/10.1101/156075>
- Douglas RJ, Martin KA. Neuronal circuits of the neocortex. *Annu Rev Neurosci*. 2004; 440:735–69.
- Druckmann S, Chklovskii DB. Neuronal circuits underlying persistent representations despite time varying activity. *Curr Biol*. 2012; 22(22):2095–103. [PubMed: 23084992]
- Durstewitz D, Seamans JK, Sejnowski TJ. Neurocomputational models of working memory. *Nat Neurosci*. 2000; 3:1184–91. [PubMed: 11127836]
- Egorov AV, Hamam BN, Fransén E, Hasselmo ME, Alonso AA. Graded persistent activity in entorhinal cortex neurons. *Nature*. 2002; 420(6912):173–78. [PubMed: 12432392]
- Faisal AA, Selen LP, Wolpert DM. Noise in the nervous system. *Nat Rev Neurosci*. 2008; 9(4):292–303. [PubMed: 18319728]
- Field GD, Sampath AP, Rieke F. Retinal processing near absolute threshold: from behavior to mechanism. *Annu Rev Physiol*. 2005; 67:491–514. [PubMed: 15709967]
- Fransén E, Tahvildari B, Egorov AV, Hasselmo ME, Alonso AA. Mechanism of graded persistent cellular activity of entorhinal cortex layer V neurons. *Neuron*. 2006; 49(5):735–46. [PubMed: 16504948]
- Fraser DD, MacVicar BA. Cholinergic-dependent plateau potential in hippocampal CA1 pyramidal neurons. *J Neurosci*. 1996; 16(13):4113–28. [PubMed: 8753873]
- Funahashi S, Bruce CJ, Goldman-Rakic PS. Mnemonic coding of visual space in the monkey’s dorsolateral prefrontal cortex. *J Neurophysiol*. 1989; 61(2):331–49. [PubMed: 2918358]
- Fusi S, Drew PJ, Abbott L. Cascade models of synaptically stored memories. *Neuron*. 2005; 45(4):599–611. [PubMed: 15721245]
- Fuster JM. Unit activity in prefrontal cortex during delayed-response performance: neuronal correlates of transient memory. *J Neurophysiol*. 1973; 36(1):61–78. [PubMed: 4196203]
- Fuster JM, Jervey JP. Inferotemporal neurons distinguish and retain behaviorally relevant features of visual stimuli. *Science*. 1981; 212:952–55. [PubMed: 7233192]
- Ganguli S, Biseley JW, Roitman JD, Shadlen MN, Goldberg ME, Miller KD. One-dimensional dynamics of attention and decision making in LIP. *Neuron*. 2008; 58(1):15–25. [PubMed: 18400159]
- Ganguli S, Latham P. Feedforward to the past: the relation between neuronal connectivity, amplification, and short-term memory. *Neuron*. 2009; 61(4):499–501. [PubMed: 19249270]
- Gnadt JW, Andersen RA. Memory related motor planning activity in posterior parietal cortex of macaque. *Exp Brain Res*. 1988; 70:216–20. [PubMed: 3402565]
- Goldman MS. Memory without feedback in a neural network. *Neuron*. 2009; 61(4):621–34. [PubMed: 19249281]
- Goldman MS, Levine JH, Major G, Tank DW, Seung H. Robust persistent neural activity in a model integrator with multiple hysteretic dendrites per neuron. *Cereb Cortex*. 2003; 13(11):1185–95. [PubMed: 14576210]
- Goldman-Rakic PS. Cellular basis of working memory. *Neuron*. 1995; 14(3):477–85. [PubMed: 7695894]

- Goldwyn JH, Shea-Brown E. The what and where of adding channel noise to the Hodgkin-Huxley equations. *PLOS Comput Biol.* 2011; 7:e1002247. [PubMed: 22125479]
- Gottlieb Y, Vaadia E, Abeles M. Single unit activity in the auditory cortex of a monkey performing a short term memory task. *Exp Brain Res.* 1989; 74:139–48. [PubMed: 2924831]
- Guinamard R, Simard C, Del Negro C. Flufenamic acid as an ion channel modulator. *Pharmacol Ther.* 2013; 138(2):272–84. [PubMed: 23356979]
- Haj-Dahmane S, Andrade R. Muscarinic activation of a voltage-dependent cation nonselective current in rat association cortex. *J Neurosci.* 1996; 16(12):3848–61. [PubMed: 8656279]
- Haj-Dahmane S, Andrade R. Muscarinic receptors regulate two different calcium-dependent non-selective cation currents in rat prefrontal cortex. *Eur J Neurosci.* 1999; 11(6):1973–80. [PubMed: 10336666]
- Harvey CD, Coen P, Tank DW. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature.* 2012; 484(7392):62–68. [PubMed: 22419153]
- Hasselmo ME, Stern CE. Mechanisms underlying working memory for novel information. *Trends Cogn Sci.* 2006; 10(11):487–93. [PubMed: 17015030]
- Hernández A, Zainos A, Romo R. Temporal evolution of a decision-making process in medial premotor cortex. *Neuron.* 2002; 33(6):959–72. [PubMed: 11906701]
- Hillar C, Sohl-Dickstein J, Koepsell K. Efficient and optimal binary hopfield associative memory storage using minimum probability flow. 2012 arXiv:1204.2916 [nlin.AO].
- Hopfield JJ. Neural networks and physical systems with emergent collective computational abilities. *PNAS.* 1982; 79(8):2554–58. [PubMed: 6953413]
- Hughes SW, Cope DW, Tóth TI, Williams SR, Crunelli V. All thalamocortical neurones possess a T-type Ca^{2+} ‘window’ current that enables the expression of bistability-mediated activities. *J Physiol.* 1999; 517(3):805–15. [PubMed: 10358120]
- Hyde RA, Strowbridge BW. Mnemonic representations of transient stimuli and temporal sequences in the rodent hippocampus in vitro. *Nat Neurosci.* 2012; 15(10):1430–38. [PubMed: 22960934]
- Jentsch TJ. Neuronal KCNQ potassium channels: physiology and role in disease. *Nat Rev Neurosci.* 2000; 1(1):21–30. [PubMed: 11252765]
- Jochems A, Yoshida M. Persistent firing supported by an intrinsic cellular mechanism in hippocampal CA3 pyramidal cells. *Eur J Neurosci.* 2013; 38(2):2250–59. [PubMed: 23651161]
- Joshua M, Lisberger S. A tale of two species: neural integration in zebrafish and monkeys. *Neuroscience.* 2015; 296:80–91. [PubMed: 24797331]
- Kass JI, Mintz IM. Silent plateau potentials, rhythmic bursts, and pacemaker firing: three patterns of activity that coexist in quadristable subthalamic neurons. *PNAS.* 2006; 103(1):183–88. [PubMed: 16373507]
- Kilpatrick ZP, Ermentrout B, Doiron B. Optimizing working memory with heterogeneity of recurrent cortical excitation. *J Neurosci.* 2013; 33:18999–9011. [PubMed: 24285904]
- Knauer B, Jochems A, Valero-Aracama MJ, Yoshida M. Long-lasting intrinsic persistent firing in rat CA1 pyramidal cells: a possible mechanism for active maintenance of memory. *Hippocampus.* 2013; 23(9):820–31. [PubMed: 23609880]
- Kopec CD, Erlich JC, Brunton BW, Deisseroth K, Brody CD. Cortical and subcortical contributions to short-term memory for orienting movements. *Neuron.* 2015; 88(2):367–77. [PubMed: 26439529]
- Koulakov AA, Raghavachari S, Kepecs A, Lisman JE. Model for a robust neural integrator. *Nat Neurosci.* 2002; 5(8):775–82. [PubMed: 12134153]
- Kubota K, Iwamoto T, Suzuki H. Visuokinetic activities of primate prefrontal neurons during delayed-response performance. *J Neurophysiol.* 1974; 37:1197–212. [PubMed: 4436696]
- Lahiri, S., Ganguli, S., Burges, CJC, Bottou, L., Welling, M., Ghahramani, Z., Weinberger, KQ., editors. A memory frontier for complex synapses; *Advances in Neural Information Processing Systems* 26. 2013. p. 1034–42. <https://papers.nips.cc/paper/4872-a-memory-frontier-for-complex-synapses>
- Laje R, Buonomano DV. Robust timing and motor patterns by taming chaos in recurrent neural networks. *Nat Neurosci.* 2013; 16(7):925–33. [PubMed: 23708144]
- Larimer P, Strowbridge BW. Nonrandom local circuits in the dentate gyrus. *J Neurosci.* 2008; 28(47):12212–23. [PubMed: 19020015]

- Larimer P, Strowbridge BW. Representing information in cell assemblies: persistent activity mediated by semilunar granule cells. *Nat Neurosci.* 2010; 13(2):213–22. [PubMed: 20037579]
- Lei YT, Thuault SJ, Launay P, Margolskee RF, Kandel ER, Siegelbaum SA. Differential contribution of TRPM4 and TRPM5 nonselective cation channels to the slow afterdepolarization in mouse prefrontal cortex neurons. *Front Cell Neurosci.* 2014; 8:267. [PubMed: 25237295]
- Li N, Daie K, Svoboda K, Druckmann S. Robust neuronal dynamics in premotor cortex during motor planning. *Nature.* 2016; 532(7600):459–64. [PubMed: 27074502]
- Lim S, Goldman MS. Balanced cortical microcircuitry for maintaining information in working memory. *Nat Neurosci.* 2013; 16(9):1306–14. [PubMed: 23955560]
- Lisman JE, Goldring MA. Feasibility of long-term storage of graded information by the Ca²⁺/calmodulin-dependent protein kinase molecules of the postsynaptic density. *PNAS.* 1988; 85(14):5320–24. [PubMed: 3393540]
- Litwin-Kumar A, Doiron B. Formation and maintenance of neuronal assemblies through synaptic plasticity. *Nat Commun.* 2014; 5:5319. [PubMed: 25395015]
- Liu D, Gu X, Zhu J, Zhang X, Han Z, et al. Medial prefrontal activity during delay period contributes to learning of a working memory task. *Science.* 2014; 346(6208):458–63. [PubMed: 25342800]
- Maass W, Natschläger T, Markram H. Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Comput.* 2002; 14(11):2531–60. [PubMed: 12433288]
- MacDonald CJ, Lepage KQ, Eden UT, Eichenbaum H. Hippocampal “time cells” bridge the gap in memory for discontinuous events. *Neuron.* 2011; 71(4):737–49. [PubMed: 21867888]
- MacNeil D, Eliasmith C. Fine-tuning and the stability of recurrent neural networks. *PLOS ONE.* 2011; 6(9):e22885. [PubMed: 21980334]
- Major G, Tank D. Persistent neural activity: prevalence and mechanisms. *Curr Opin Neurobiol.* 2004; 14(6):675–84. [PubMed: 15582368]
- Miles R, Wong RK. Excitatory synaptic interactions between CA3 neurones in the guinea-pig hippocampus. *J Physiol.* 1986; 373:397–418. [PubMed: 3018233]
- Miles R, Wong RK. Latent synaptic pathways revealed after tetanic stimulation in the hippocampus. *Nature.* 1987; 329(6141):724–26. [PubMed: 3670374]
- Miller EK, Erickson CA, Desimone R. Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *J Neurosci.* 1996; 16:5154–67. [PubMed: 8756444]
- Miller EK, Li L, Desimone R. Activity of neurons in anterior inferior temporal cortex during a short-term memory task. *J Neurosci.* 1993; 13:1460–78. [PubMed: 8463829]
- Miyashita Y, Chang HS. Neuronal correlate of pictorial short-term memory in the primate temporal cortex. *Nature.* 1988; 331:68–70. [PubMed: 3340148]
- Mongillo G, Barak O, Tsodyks M. Synaptic theory of working memory. *Science.* 2008; 319(5869):1543–46. [PubMed: 18339943]
- Morishima M, Morita K, Kubota Y, Kawaguchi Y. Highly differentiated projection-specific cortical subnetworks. *J Neurosci.* 2011; 31(28):10380–91. [PubMed: 21753015]
- Murray JD, Bernacchia A, Roy NA, Constantinidis C, Romo R, Wang X-J. Stable population coding for working memory coexists with heterogeneous neural dynamics in prefrontal cortex. *PNAS.* 2016; 114(2):394–99. [PubMed: 28028221]
- Nastainczyk W, Meves H, Watt DD. A short-chain peptide toxin isolated from *Centruroides sculpturatus* scorpion venom inhibits ether-à-go-go-related gene K⁺ channels. *Toxicol.* 2002; 40(7):1053–58. [PubMed: 12076661]
- Navaroli VL, Zhao Y, Boguszewski P, Brown TH. Muscarinic receptor activation enables persistent firing in pyramidal neurons from superficial layers of dorsal perirhinal cortex. *Hippocampus.* 2012; 22(6):1392–404. [PubMed: 21956787]
- Niessing J, Friedrich RW. Olfactory pattern classification by discrete neuronal network states. *Nature.* 2010; 465(7294):47–52. [PubMed: 20393466]
- Otsuka T, Murakami F, Song WJ. Excitatory postsynaptic potentials trigger a plateau potential in rat subthalamic neurons at hyperpolarized states. *J Neurophysiol.* 2001; 86(4):1816–25. [PubMed: 11600642]

- Pace RW, Mackay DD, Feldman JL, Del Negro CA. Inspiratory bursts in the preBötzing complex depend on a calcium-activated non-specific cation current linked to glutamate receptors in neonatal mice. *J Physiol.* 2007; 582(1):113–25. [PubMed: 17446214]
- Panzeri S, Harvey CD, Piasini E, Latham PE, Fellin T. Cracking the neural code for sensory perception by combining statistics, intervention, and behavior. *Neuron.* 2017; 93:491–507. [PubMed: 28182905]
- Pastalkova E, Itskov V, Amarasingham A, Buzsaki G. Internally generated cell assembly sequences in the rat hippocampus. *Science.* 2008; 321(5894):1322–27. [PubMed: 18772431]
- Perrier JF, Alaburda A, Hounsgaard J. Spinal plasticity mediated by postsynaptic L-type Ca^{2+} channels. *Brain Res Brain Res Rev.* 2002; 40(1–3):223–29. [PubMed: 12589920]
- Pesaran B, Pezaris JS, Sahani M, Mitra PP, Andersen RA. Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nat Neurosci.* 2002; 5:805–11. [PubMed: 12134152]
- Pessia M, Servetini I, Panichi R, Guasti L, Grassi S, et al. ERG voltage-gated K^{+} channels regulate excitability and discharge dynamics of the medial vestibular nucleus neurones. *J Physiol.* 2008; 586(20):4877–90. [PubMed: 18718985]
- Petersson ME, Fransén E. Long-lasting small-amplitude TRP-mediated dendritic depolarizations in CA1 pyramidal neurons are intrinsically stable and originate from distal tuft regions. *Eur J Neurosci.* 2012; 36(7):2917–25. [PubMed: 22758919]
- Pinsky PF, Rinzel J. Intrinsic and network rhythmogenesis in a reduced Traub model for CA3 neurons. *J Comput Neurosci.* 1994; 1(1–2):39–60. [PubMed: 8792224]
- Pressler RT, Strowbridge BW. Blanes cells mediate persistent feedforward inhibition onto granule cells in the olfactory bulb. *Neuron.* 2006; 49(6):889–904. [PubMed: 16543136]
- Rahman J, Berger T. Persistent activity in layer 5 pyramidal neurons following cholinergic activation of mouse primary cortices. *Eur J Neurosci.* 2011; 34(1):22–30. [PubMed: 21645136]
- Rainer G, Asaad WF, Miller EK. Selective representation of relevant information by neurons in the primate prefrontal cortex. *Nature.* 1998; 393:577–79. [PubMed: 9634233]
- Rajan K, Harvey CD, Tank DW. Recurrent network models of sequence generation and memory. *Neuron.* 2016; 90(1):128–42. [PubMed: 26971945]
- Rao SC, Rainer G, Miller EK. Integration of what and where in the primate prefrontal cortex. *Science.* 1997; 276:821–24. [PubMed: 9115211]
- Reboreda A, Jiménez-Díaz L, Navarro-López JD. TRP channels and neural persistent activity. *Adv Exp Med Biol.* 2011; 704:595–613. [PubMed: 21290318]
- Renart A, Song P, Wang X-J. Robust spatial working memory through homeostatic synaptic scaling in heterogeneous cortical networks. *Neuron.* 2003; 38(3):473–85. [PubMed: 12741993]
- Rigotti M, Barak O, Warden MR, Wang X-J, Daw ND, et al. The importance of mixed selectivity in complex cognitive tasks. *Nature.* 2013; 497(7451):585–90. [PubMed: 23685452]
- Romo R, Brody CD, Hernández A, Lemus L. Neuronal correlates of parametric working memory in the prefrontal cortex. *Nature.* 1999; 399(6735):470–73. [PubMed: 10365959]
- Rubin JE, Hayes JA, Mendenhall JL, Del Negro CA. Calcium-activated nonspecific cation current and synaptic depression promote network-dependent burst oscillations. *PNAS.* 2009; 106(8):2939–44. [PubMed: 19196976]
- Russo RE, Hounsgaard J. Plateau-generating neurones in the dorsal horn in an in vitro preparation of the turtle spinal cord. *J Physiol.* 1996; 493(1):39–54. [PubMed: 8735693]
- Santos MD, Mohammadi MH, Yang S, Liang CW, Kao JP, et al. Dendritic hold and read: a gated mechanism for short term information storage and retrieval. *PLOS ONE.* 2012; 7(5):e37542. [PubMed: 22629416]
- Savin C, Triesch J. Emergence of task-dependent representations in working memory circuits. *Front Comput Neurosci.* 2014; 8:57. [PubMed: 24904395]
- Schwemmer MA, Fairhall AL, Denéve S, Shea-Brown ET. Constructing precisely computing networks with biophysical spiking neurons. *J Neurosci.* 2015; 35(28):10112–34. [PubMed: 26180189]
- Seung HS. How the brain keeps the eyes still. *PNAS.* 1996; 93(23):13339–44. [PubMed: 8917592]

- Shafi M, Zhou Y, Quintana J, Chow C, Fuster J, Bodner M. Variability in neuronal activity in primate cortex during working memory tasks. *Neuroscience*. 2007; 146(3):1082–108. [PubMed: 17418956]
- Singh R, Eliasmith C. Higher-dimensional neurons explain the tuning and dynamics of working memory cells. *J Neurosci*. 2006; 26(14):3667–78. [PubMed: 16597721]
- Super H, Spekreijse H, Lamme VA. A neural correlate of working memory in the monkey primary visual cortex. *Science*. 2001; 293:120–124. [PubMed: 11441187]
- Tahvildari B, Alonso AA, Bourque CW. Ionic basis of ON and OFF persistent activity in layer III lateral entorhinal cortical principal neurons. *J Neurophysiol*. 2008; 99(4):2006–11. [PubMed: 18256167]
- Tegnér J, Compte A, Wang X-J. The dynamical stability of reverberatory neural circuits. *Biol Cybern*. 2002; 87:471–481. [PubMed: 12461636]
- Toyoizumi T, Abbott L. Beyond the edge of chaos: amplification and temporal integration by recurrent networks in the chaotic regime. *Phys Rev E*. 2011; 84(5):051908.
- Traub RD, Jefferys JGR. Are there unifying principles underlying the generation of epileptic afterdischarges in vitro? *Prog Brain Res*. 1994; 102:383–94. [PubMed: 7800828]
- Traub RD, Wong RK. Cellular mechanism of neuronal synchronization in epilepsy. *Science*. 1982; 216(4547):745–47. [PubMed: 7079735]
- Trengove C, van Leeuwen C, Diesmann M. High-capacity embedding of synfire chains in a cortical network model. *J Comput Neurosci*. 2013; 34(2):185–209. [PubMed: 22878688]
- Vergara J, Rivera N, Rossi-Pool R, Romo R. A neural parametric code for storing information of more than one sensory modality in working memory. *Neuron*. 2016; 89(1):54–62. [PubMed: 26711117]
- Wang X-J. Synaptic basis of cortical persistent activity: the importance of NMDA receptors to working memory. *J Neurosci*. 1999; 19:9587–603. [PubMed: 10531461]
- Wang X-J. Synaptic reverberation underlying mnemonic persistent activity. *Trends Neurosci*. 2001; 24(8):455–63. [PubMed: 11476885]
- White JA, Rubinstein JT, Kay AR. Channel noise in neurons. *Trends Neurosci*. 2000; 23:131–37. [PubMed: 10675918]
- Williams PA, Larimer P, Gao Y, Strowbridge BW. Semilunar granule cells: glutamatergic neurons in the rat dentate gyrus with axon collaterals in the inner molecular layer. *J Neurosci*. 2007; 27(50):13756–61. [PubMed: 18077687]
- Wills TJ, Lever C, Cacucci F, Burgess N, O’Keefe J. Attractor dynamics in the hippocampal representation of the local environment. *Science*. 2005; 308:873–76. [PubMed: 15879220]
- Wimmer K, Nykamp DQ, Constantinidis C, Compte A. Bump attractor dynamics in prefrontal cortex explains behavioral precision in spatial working memory. *Nat Neurosci*. 2014; 17(3):431–39. [PubMed: 24487232]
- Wong RK, Traub RD, Miles R. Cellular basis of neuronal synchrony in epilepsy. *Adv Neurol*. 1986; 44:583–92. [PubMed: 3706021]
- Yamada-Hanff J, Bean BP. Persistent sodium current drives conditional pacemaking in CA1 pyramidal neurons under muscarinic stimulation. *J Neurosci*. 2013; 33(38):15011–21. [PubMed: 24048831]
- Yan HD, Villalobos C, Andrade R. TRPC channels mediate a muscarinic receptor-induced afterdepolarization in cerebral cortex. *J Neurosci*. 2009; 29(32):10038–46. [PubMed: 19675237]
- Yoon K, Buice MA, Barry C, Hayman R, Burgess N, Fiete IR. Specific evidence of low-dimensional continuous attractor dynamics in grid cells. *Nat Neurosci*. 2013; 16:1077–84. [PubMed: 23852111]
- Zaksas D, Pasternak T. Directional signals in the prefrontal cortex and in area MT during a working memory for visual motion task. *J Neurosci*. 2006; 26:11726–42. [PubMed: 17093094]
- Zenke F, Agnes EJ, Gerstner W. Diverse synaptic plasticity mechanisms orchestrated to form and retrieve memories in spiking neural networks. *Nat Commun*. 2015; 6:6922. [PubMed: 25897632]
- Zhang Z, Reboreda A, Alonso A, Barker PA, Séguéla P. TRPC channels underlie cholinergic plateau potentials and persistent activity in entorhinal cortex. *Hippocampus*. 2011; 21(4):386–97. [PubMed: 20082292]

- Zhou YD, Fuster JM. Mnemonic neuronal activity in somatosensory cortex. *PNAS*. 1996; 93:10533–37. [PubMed: 8927629]
- Zylberberg J, Cafaro J, Turner MH, Shea-Brown E, Rieke F. Direction-selective circuits shape noise to ensure a precise population code. *Neuron*. 2016a; 89:369–83. [PubMed: 26796691]
- Zylberberg J, Hyde RA, Strowbridge BW. Dynamics of robust pattern separability in the hippocampal dentate gyrus. *Hippocampus*. 2016b; 26(5):623–32. [PubMed: 26482936]

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

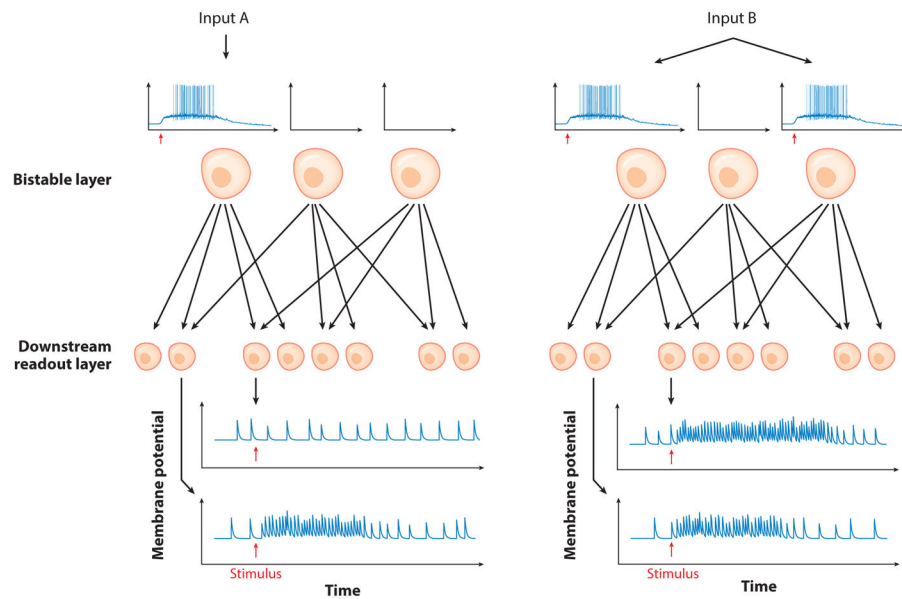


Figure 1.

Memory function created by embedding bistable neurons in feedforward networks. A network is modeled on the dentate gyrus, which contains a set of bistable excitatory neurons (semilunar granule cells) (Larimer & Strowbridge 2010, Williams et al. 2007) that generate divergent projections onto downstream cells (hilar neurons in the dentate gyrus). Different input patterns can trigger bistable responses in different subsets of bistable cells, leading to different combinations of synaptic input barrages in the downstream neurons. Which input was presented can be determined by monitoring changes in the synaptic tone in a small subgroup of downstream neurons (as in Hyde & Strowbridge 2012 and Larimer & Strowbridge 2010).

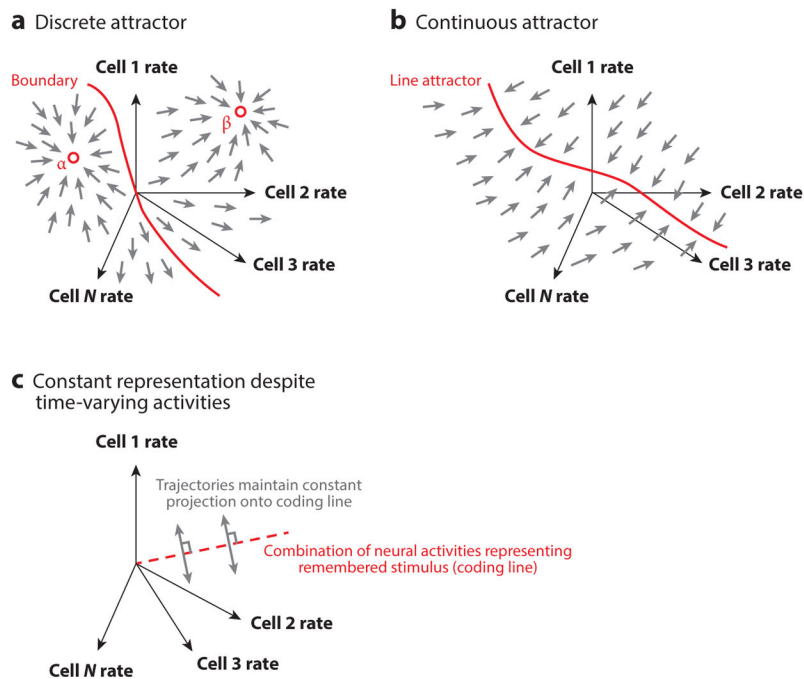


Figure 2.

Network models capable of generating persistent representations. (a) In discrete attractor models, a population of neurons is described by their firing rates (axes of the diagram). The network dynamics cause movement within this space: At each point, the small arrows indicate the direction in which the population activities move. (This is known as a direction field in mathematics, and one can visualize trajectories of the population by connecting neighboring arrows in a head-to-tail fashion.) Here, all direction arrows point toward either point α or point β , and so the network activity patterns will evolve toward one of these two activity patterns. Which of these patterns gets generated depends on whether the inputs push the network to the left of the marked boundary or the right. This boundary is known as a separatrix. (b) Continuous attractor models are similar to the model in panel a, but now the direction arrows all point toward a continuous line. The network dynamics cause the activity patterns to evolve to points (patterns) on the marked line. (c) In models displaying continuous representations despite time-varying neural activities, the remembered stimulus value is assumed to be encoded in a combination of neural firing rates: in other words, by the projection of the population firing rate vector onto a line (coding line) in the space of neural activities. Here, the dynamical evolution of the neural activities is orthogonal to that coding line, and so the changes in neural firing rates do not change the projection of the firing rate vector onto that line. Consequently, the representation is stably maintained.