

Beyond working memory: the role of persistent activity in decision making

Clayton E. Curtis¹ and Daeyeol Lee²

¹ Department of Psychology and Center for Neural Science, New York University, 6 Washington Place, New York, NY 10003, USA

² Department of Neurobiology, Yale University School of Medicine, 333 Cedar Street, New Haven, CT 06510, USA

Since its first discovery in the prefrontal cortex, persistent activity during the interval between a transient sensory stimulus and a subsequent behavioral response has been identified in many cortical and subcortical areas. Such persistent activity is thought to reflect the maintenance of working memory representations that bridge past events with future contingent plans. Indeed, the term persistent activity is sometimes used interchangeably with working memory. In this review, we argue that persistent activity observed broadly across many cortical and subcortical areas reflects not only working memory maintenance, but also a variety of other cognitive processes, including perceptual and reward-based decision making.

Properties of persistent activity

For decades now we have known that neuronal activity persists in the prefrontal cortex when the subject is in an active state of remembering a recent event [1–3], a process termed working memory [4,5]. Working memory allows animals to use information that is not currently available in the environment but is crucial for adaptive behavior. In other words, an internal representation of relevant information must be created and maintained until it can be used later to guide behavior.

Typical working memory experiments involve the brief presentation of a stimulus to be remembered, which is then followed by a short delay, or retention interval. Performance is probed by having the subject make a response that is contingent on successful maintenance of the memoranda. Neural activity in the prefrontal cortex persists during the delay period without being contaminated by sensory and motor processes, so it is thought to reflect the maintenance of a working memory representation [1–3].

Several key features of persistent activity suggest that it is a neural mechanism critical for working memory maintenance and presumably other high-level cognitive processes such as decision making. First, neural activity, measured by a variety of methodologies, such as single-neuron recording and neuroimaging, persists during the time epoch when a representation is thought to be maintained in an active state (e.g. when a stimulus needs to be remembered) [1–3,5–7]. Second, sustained neural activity dissipates when the representation is no longer needed (e.g. when a memory-guided response has been generated) [3,8,9]. Third, when activity does not persist throughout a

retention interval, memory performance is compromised [3,10]. Fourth, persistent activity scales with working memory load. For example, as the number of items that need to be maintained increases, the level of delay period activity also increases until it plateaus near the limits of the individual's short-term memory capacity [11–13]. Fifth, persistent activity is selective, such that specific information about the contents of working memory can be decoded from the population activity [14–16]. For example, persistent activity is selective for the location of the memoranda during spatial working memory tasks [3,9] (Figure 1). Finally, although much of the relevant research has focused on the prefrontal cortex, persistent activity has been reported in many other cortical and subcortical areas [17,18], suggesting that its generation and maintenance might depend on a broadly distributed circuit mechanism (Box 1). In this review, we argue that the persistence of neural activity in the absence of external stimulation is the general mechanism by which internal representations are maintained in an active state. Moreover, these representations go beyond traditional notions of working memory and include variables useful for decision making [19,20].

Persistent activity and working memory

The development of neurobiological theories of working memory critically depends on the measurement of neural activity during the delay period. As a result, persistent activity has essentially become a proxy for working memory and the two are often used interchangeably. This is not surprising because persistent neural activity during delay periods seems to carry useful information about the features of the memoranda. Persistent activity arises while the subject maintains information about the spatial location [3,9,10,17,18,21,22], object identity [23–25], word [26], sound [27,28] or haptic sensation [29] of the memorandum. Theoretically, the short-term maintenance of such features is necessary for a variety of high-level cognitive functions because it facilitates temporal integration of past sensory events with future actions [5,30]. Importantly, working memory paradigms have been the test bed for studies addressing the neural mechanisms of persistent activity (Box 1).

Despite its logical fallibility, the reverse inference that the presence of persistent activity indicates the maintenance of a past stimulus continues to tempt researchers. The problem, however, is that persistent activity has been linked to several other cognitive processes in addition to

Corresponding author: Lee, D. (daeyeol.lee@yale.edu).

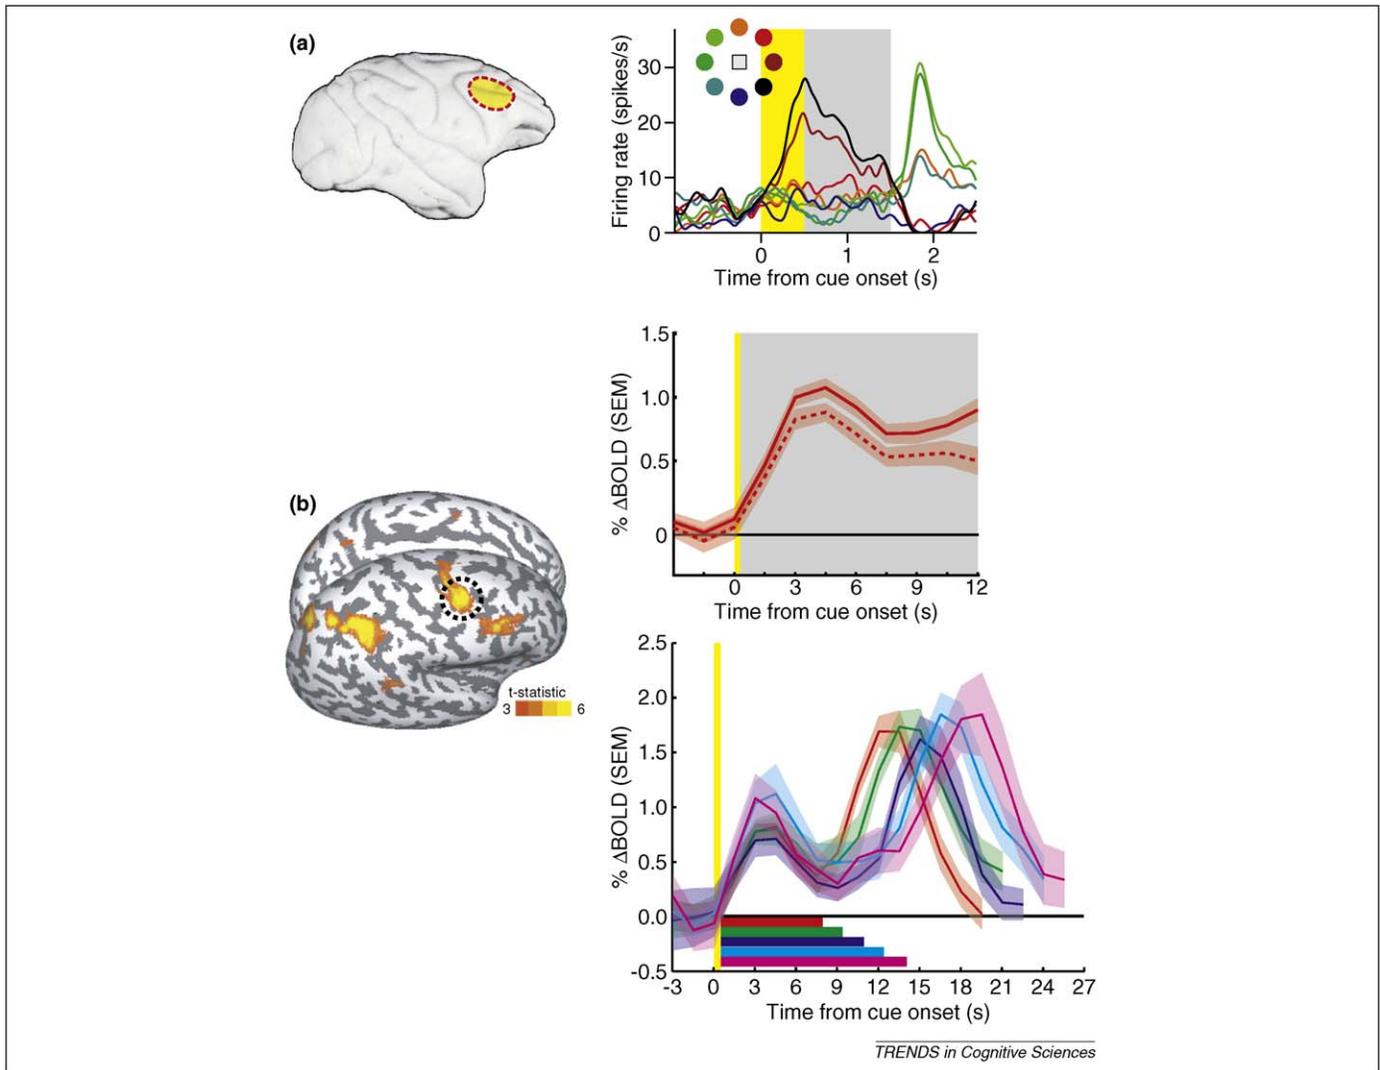


Figure 1. Persistent activity for spatial working memory. **(a)** Left, lateral view of the rhesus monkey brain. Right, spike density functions of a single neuron in the dorsolateral prefrontal cortex of a rhesus monkey during the delay period (gray background), with colors indicating the remembered location of the visual stimulus (see inset) briefly flashed during the cue period (yellow background). During the experiment, all peripheral stimuli were presented in the same color [20,48]. **(b)** Left, maintenance of a spatial location during working memory task- evoked BOLD activity in the frontal and parietal cortices, shown as a statistical map overlaid on the inflated cortex (sulci, dark gray; gyri, white). Top right, a BOLD signal persisted during a spatial working memory task in the dorsolateral prefrontal cortex (dotted circle in the left panel) and was greater for memoranda in the contralateral visual field (solid line) than for those in the ipsilateral field (dashed line). The yellow bar represents presentation of the sample cue and the gray background depicts the memory delay. Bottom right, time course of BOLD signals from the dorsolateral prefrontal cortex aligned on presentation of the sample cue during a spatial working memory task. Separate lines represent the different delay lengths (indicated by colored bars at the bottom). Importantly, persistent activity bracketed by the phasic BOLD signals increased in duration with the delay length and was sustained until the working memory representation could be used to guide the response [10].

working memory [6,31,32]. Furthermore, neural activity in the prefrontal cortex still persists and encodes the properties of visual stimuli even in the absence of working memory demands [33]. Persistent activity can also reflect not the maintenance of past events, but the anticipation of future ones. During an interval prior to or in preparation for an action, persistent activity is thought to represent the metrics or goal of a forthcoming planned action [34,35]. Similarly, persistent activity has been linked to stimuli that are anticipated but not yet encountered [25,36]. Therefore, persistent activity underlies the representation of retrospective, current and prospective information. Moreover, persistent activity can be observed in the prefrontal cortex during the active representation of more abstract information, such as rules, associations, categories and strategies [37–39]. Therefore, in addition to traditional notions of working memory representations,

persistent activity can reflect a wider variety of cognitive processes, some of which are used for making decisions and to which we now turn our attention.

Persistent activity reflecting accumulation of sensory evidence

As described above, activity of individual neurons and blood-oxygen-level-dependent (BOLD) signals in multiple brain areas can be sustained during the delay period of a working memory task, and therefore might provide the substrate for working memory and storage of other behaviorally relevant information. It is often thought that such persistent activity might arise from a network of recurrently connected neurons [19]. However, persistent activity related to a past sensory event might also arise indirectly from the short-term changes in synaptic efficacy induced by transient neural activity [40,41] (Box 1). Regardless of its

Box 1. Mechanisms of persistent activity

Precisely how persistent activity is generated and controlled in the mammalian brain is not known, but several possible, and not mutually exclusive, mechanisms have been proposed. Given the ubiquity of persistent activity, it remains an important area of future research to determine the exact relationship between each of these mechanisms and a specific form of persistent activity observed during a particular behavioral task.

- Persistent activity from reverberations in a recurrent network. Persistent activity might be initiated and sustained by reciprocal positive feedback within a population of neurons. Persistent activity observed in a given cortical area (e.g. prefrontal cortex) might arise locally within a population of neighboring neurons [19], within a population of neurons distributed across multiple cortical regions [21] or from a network of areas including both cortical and subcortical systems, such as the corticobasal ganglia–thalamus–cortical loops [17,18]. At each of these network levels, activation of separate networks could specify different stimuli, different actions or stimulus–response mappings, different behavioral strategies and reward expectancies.
- Persistent activity relying on intracellular signals. Individual cortical neurons might display persistent activity due to specific membrane currents or cumulative changes in the concentration of intracellular calcium. For example, neurons in the entorhinal cortex can display a step-like depolarization following brief stimulation. Increased spiking can lead to calcium influx through voltage-gated calcium channels at dendrites, causing subsequent depolarization and thereby closing a positive feedback loop. This form of persistent activity seems to rely on nonspecific calcium-sensitive cationic currents [67].
- Persistent activity and short-term synaptic plasticity. When a particular synapse is repeatedly stimulated, the amount of neurotransmitter released in response to a given presynaptic action potential can increase in some synapses and decrease in others [68]. The pattern of activity sustained in a network is a function of synaptic weights among its constituent neurons, so short-term synaptic changes induced by a particular pattern of activity might make it possible for the network to resume specific patterns of persistent activity even after a quiescent period [40,41].

mechanism, such persistent activity observed during a working memory task can result from temporal integration of a transient input (Figure 2a). Thus, the neurons or networks displaying persistent activity during a working memory task might be involved in other cognitive processes that require temporal integration [19,42,43]. For example, neurons displaying persistent activity might be used to average out sensory noise during perceptual decision making (Figure 2b). According to the so-called random-walk models of decision making, a behavioral response is triggered when accumulated evidence reaches a certain threshold and therefore it takes more time to identify a weak noisy sensory stimulus than a salient stimulus. Consistent with this model, during a choice reaction time task, both the speed and accuracy of responses decrease as the stimuli become more difficult to discriminate [44]. Moreover, the same neurons in the prefrontal cortex and posterior parietal cortex that show persistent activity during a working memory task might contribute to the integration of sensory evidence because their activity tends to increase more slowly during the reaction time following presentation of a relatively weak stimulus [42,45]. Similarly, fMRI studies on perceptual decision-making tasks have also found an increase in BOLD activity in the prefrontal cortex and posterior parietal cortex when the stimulus can be more easily discriminated [46].

Persistent activity in reinforcement learning

In most neurophysiological experiments conducted in behaving animals, neural activity is recorded while the animal performs a particular behavioral task repeatedly in a large number of trials separated by short inter-trial intervals. As a result, properties of persistent activity have been largely studied in relation to behaviorally relevant events within a single trial. For example, persistent activity during the delay period of a working memory task has been viewed as the substrate of working memory for the items presented at the beginning of a trial [3,8–10]. Similarly, gradual changes in neural activity observed during a perceptual decision-making task have been linked to the process of evidence accumulation within a trial [19,42,45,46]. However, the network of neurons capable of temporal integration might be exploited in tasks in which the animal is required to integrate information about the events in multiple trials. Recently, this possibility has been tested during several behavioral tasks in which animals need to change their decision-making strategies based on their previous experience [20,47–51].

The process of decision making, namely, selecting a particular action among a number of alternative actions, is trivial if the outcomes of all the alternative actions are fully known and can be evaluated along a single dimension. Often, however, the outcomes of actions vary along multiple dimensions, such as the magnitude, probability and immediacy of reward expected from each action, and decision makers need to resolve potential conflicts to maximize the magnitude and probability of their reward and to obtain it as soon as possible. Nevertheless, as long as all the information about the expected outcomes is available, such conflicts can be resolved without any additional inputs from the environment. However, it is a much more challenging task to discover an optimal decision-making strategy in an unfamiliar and stochastic environment for which the probabilities of different outcomes from alternative actions are not known. In such cases, the desirability of each action needs to be estimated through experience. The framework of reinforcement learning theory [52] parsimoniously accounts for dynamic changes in the behaviors of humans and animals as they experience unpredictable outcomes [53–55]. In addition, the same framework has played an important role in elucidating cortical and subcortical mechanisms involved in evaluating the action outcomes and exploiting this information to optimize decision-making strategies [47–49,54–58].

In reinforcement learning theory, the entire history of previous actions and their outcomes is abstracted and represented by a set of quantities, referred to as value functions. The value function (Q in Figure 2c) is an estimate of future rewards expected from a particular state and action, and is continuously adjusted according to the animal's experience. Accordingly, decision makers can make choices at any moment simply by selecting an action with the maximum value function. Value functions are increased (decreased) when the outcome from a given action is better (worse) than expected. The discrepancy between the expected and actual outcomes is referred to as the reward prediction error and might be encoded by the activity of midbrain dopamine neurons [56]. In addition, neurons with

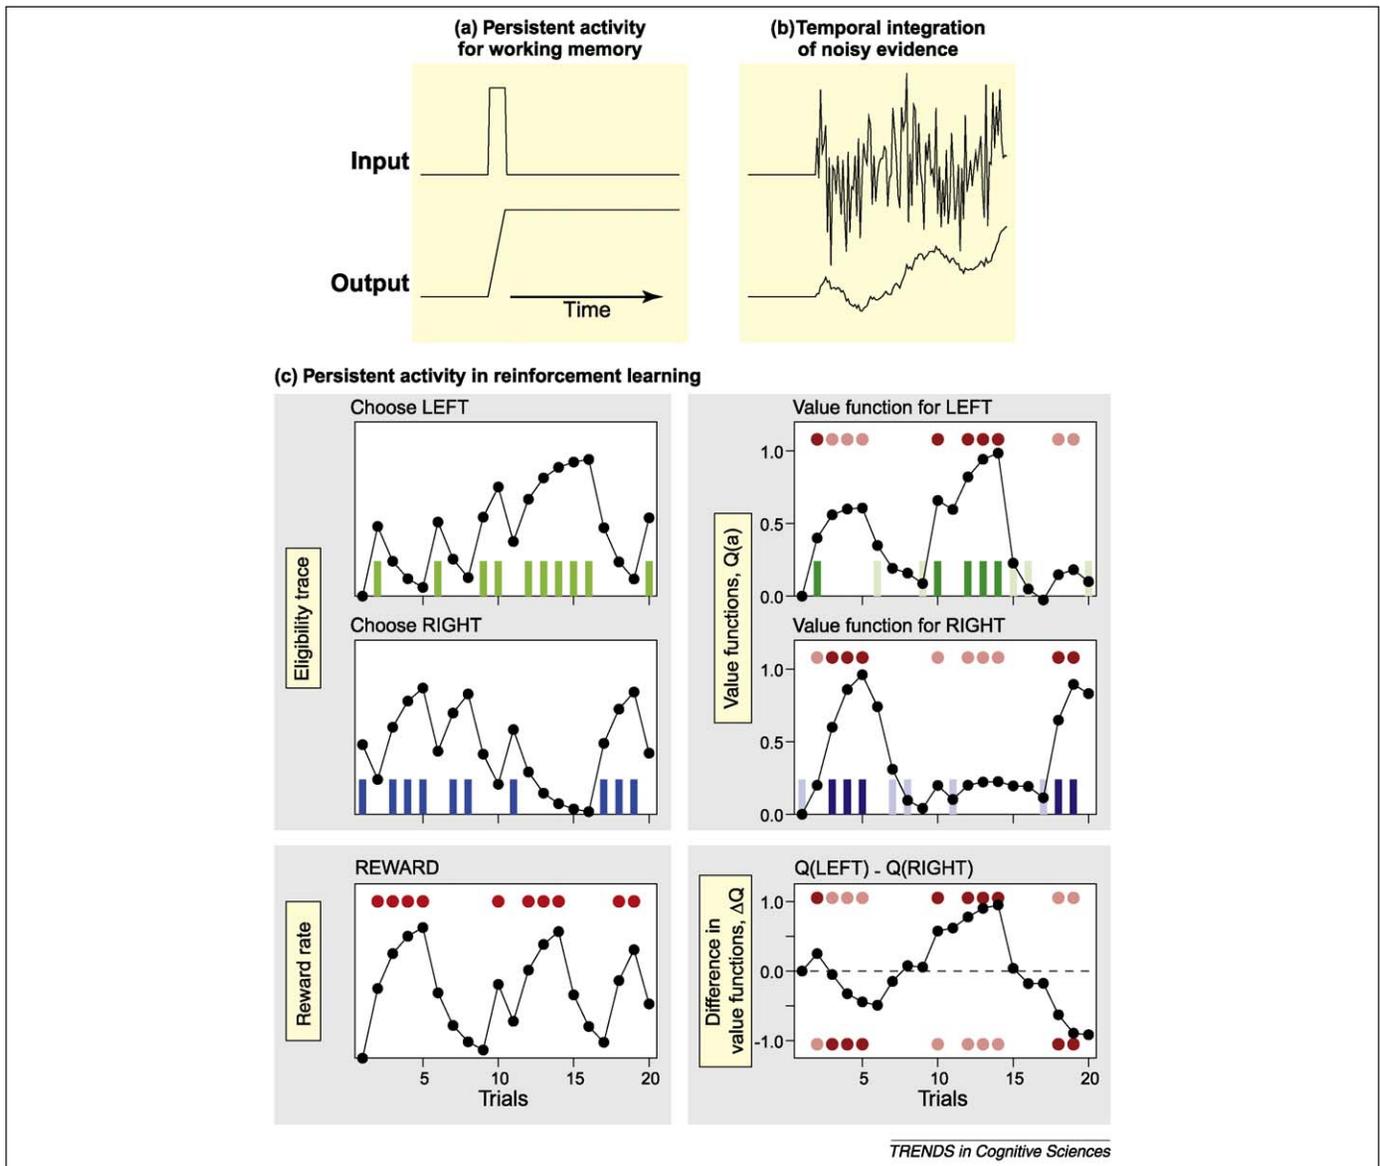


Figure 2. Persistent activity resulting from temporal integration can subserve multiple types of cognitive processes, including working memory, accumulation of noisy sensory evidence and creation of eligibility trace and other computations in reinforcement learning. (a,b) Schematic illustration of (a) the time course of inputs leading to persistent activity during a working memory task and (b) buildup activity related to the accumulation of noisy sensory inputs. (c) Hypothetical signals related to the animal's actions (green and blue vertical bars for leftward and rightward actions, respectively) can be temporally integrated to generate eligibility traces (upper and middle left panels), whereas temporal integration of reward (red disks) can lead to signals related to average reward rate (bottom left panel). The value function for each action is incremented by the reward prediction error, namely, the discrepancy between the reward earned and the reward predicted from the current value function, weighted by its eligibility trace. As a result, reward delivered in a given trial increases the value functions according to the recency of each action (upper and middle right panels), even when it was not chosen in the last trial. Dark red disks and colored vertical bars indicate the trials in which the leftward or rightward action led to immediate reward. The difference in value functions for the two alternative actions is also shown [bottom right panel; dark red circles indicate the trials in which reward was obtained after choosing the leftward (top) and rightward (bottom) actions, respectively].

a change in activity according to the value functions of specific actions have been identified in a number of brain areas, including the prefrontal cortex, posterior parietal cortex and basal ganglia [47–49,54,59–61].

As in working-memory and perceptual decision-making tasks, persistent activity has been observed in multiple brain areas during reward-based decision-making tasks. Activity related to value functions is often manifest as persistent activity [43,47–49,51,61] and therefore can be used to guide the selection of subsequent actions. Moreover, many of the key computations in reinforcement learning correspond to temporal integration and might be implemented in the form of persistent activity. For example, the outcomes from a particular action are often

revealed to the animal after substantial delays. In some cases, the animal might execute multiple actions before receiving a reward or penalty. This makes it difficult to determine how the value functions should be adjusted for multiple actions, creating a problem of temporal credit assignment [52,62]. In the so-called temporal-difference (TD) learning algorithms, this problem is addressed by comparing the value functions of two successive states. Therefore, a particular action can produce a positive reward prediction error and increase its value function, even without receiving any reward immediately, if it allows the decision maker to move to a state with a higher value function than that for the current state. Although TD learning provides a solution to the problem of temporal

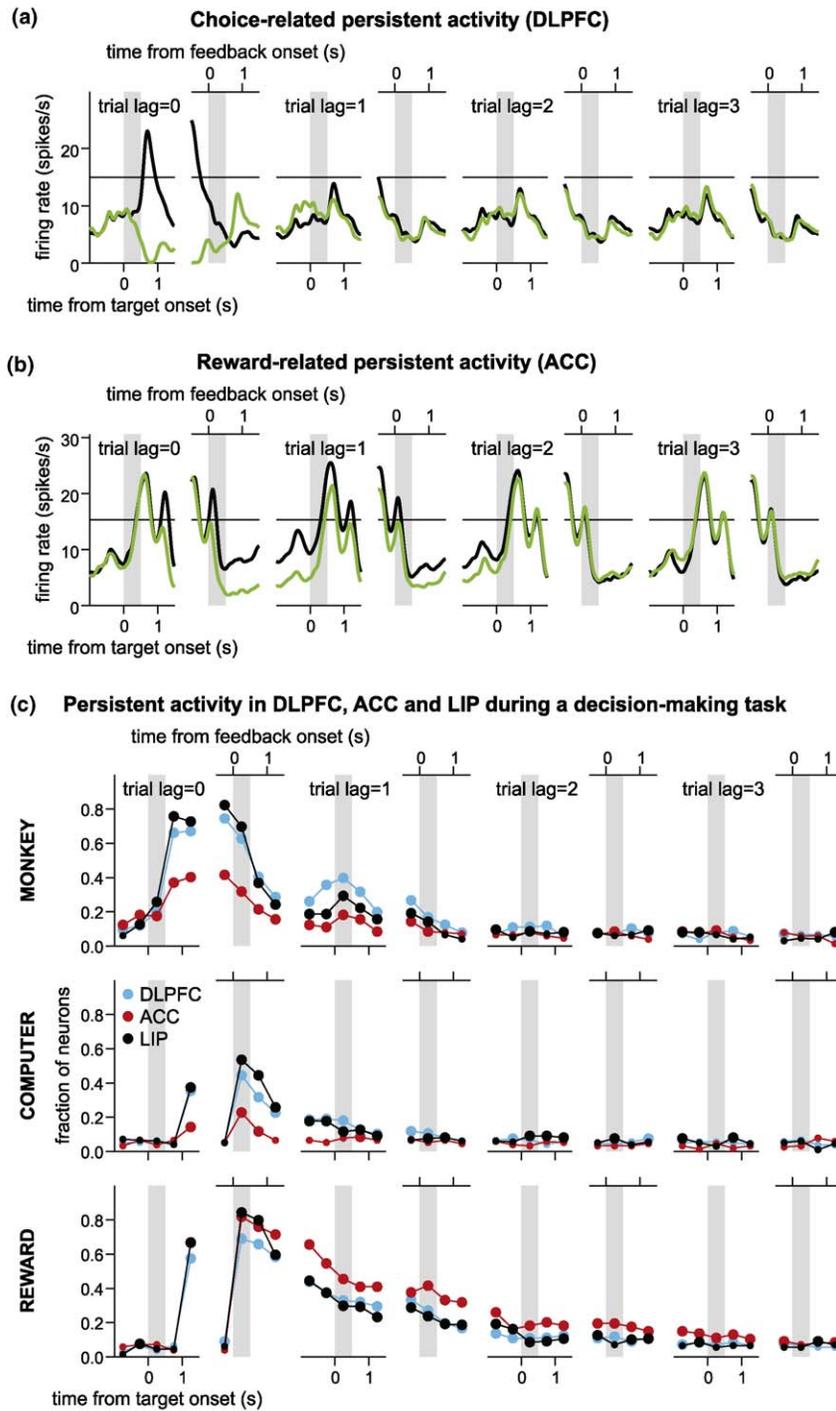


Figure 3. Persistent activity for decision making and reinforcement learning. **(a)** Persistent activity related to the animal's choice [20]. The spike density function of the same neuron illustrated in Figure 1a during a binary decision-making task is plotted separately according to the position of the target chosen in the current trial (trial lag=0) or in each of the previous three trials (trial lag=1–3). Black and green lines indicate the activity for the trials in which the animal chose the leftward and rightward targets, respectively. **(b)** Persistent activity related to the outcome of the animal's decision [43]. The spike density functions of a neuron in the anterior cingulate cortex is plotted separately according to whether the animal was rewarded (green) or not (black) in the current (trial lag=0) or in each of the previous three trials (trial lag=1–3) using the same format as in **(a)**. **(c)** Proportion of neurons in the dorsolateral prefrontal cortex (DLPFC), anterior cingulate cortex (ACC) and posterior parietal cortex (lateral intra-parietal area, LIP) that showed a significant change in activity according to the animal's choice (top), the choice of the computer opponent (middle), and outcome of the animal's choice (bottom) in the current (trial lag=0) and previous three trials (trial lag=1–3). Two subpanels for each trial lag illustrate the activity aligned at the target onset (left) and feedback onset (right), and gray bars correspond to the 0.5-s cue period and feedback period, respectively. The decision-making task used in these studies simulated a matching-pennies task in which the animal was rewarded only when it chose the same target as the computer opponent [20,43,47–49,53].

credit assignment, the time course of reward prediction errors encoded by dopamine neurons does not match the pattern predicted by simple TD learning algorithms, suggesting that additional mechanisms for temporal credit

assignment exist [63]. Indeed, some reinforcement learning algorithms rely on the memory of previous actions and states visited by the decision makers. These memory signals are referred to as eligibility traces [52] (Figure 2c).

Persistent activity in the prefrontal cortex [20,47,49], posterior parietal cortex [48] and basal ganglia [51,64,65] can be modulated by previous actions of the animal across multiple trials and might correspond to the eligibility traces (Figure 3a,c).

Persistent activity related to the outcomes of previous choices has also been observed in the same regions of the cortex and basal ganglia that showed persistent changes in neural activity according to the animal's previous choices [20,43,47–50,64,65] (Figures 2c and 3b,c). These results suggest that the signals related to previous actions and their outcomes might influence the ways in which sensory stimuli are encoded and transformed to behavioral responses in multiple regions of the brain. Similar to the activity related to previous choices, persistent activity related to previous outcomes often continues during inter-trial intervals. In addition, in contrast to persistent activity observed during the delay period of a working memory task, persistent activity related to the animal's previous choices and outcomes tends to decay gradually [43,48,51]. Reward-related persistent activity might provide the substrate for temporal integration of rewards in multiple trials, and could therefore be used to encode the average reward rate [43]. Information about the average reward rate plays an important role in some reinforcement learning algorithms. For example, it might be used to calculate the reward prediction error and influence how the value functions would be updated, as in so-called R-learning [52]. Indeed, the primate anterior cingulate cortex contains neurons encoding reward prediction errors in addition to neurons with persistent activity related to previous rewards [43,66]. A sudden change in the average rate is likely to indicate unexpected changes in the animal's environment that require appropriate adjustment in the animal's behavioral strategies [55]. Therefore, persistent activity related to previous reward might provide a substrate for monitoring the productivity of the animal's behavior and therefore facilitate flexible behavior.

Concluding remarks

The brain uses the information obtained from previous experience to make predictions about the future and to select behaviors that are likely to result in the most beneficial outcomes. The length of the retention interval during which the brain has to store information about past events varies greatly according to the properties of the environment and the animal's behavioral needs. Presumably, structural changes in the nervous system are necessary for long-term information storage. By contrast, persistent neural activity might underlie short-term storage of behaviorally relevant information. Such persistent activity can support diverse cognitive processes, ranging from working memory maintenance to computation of value functions. An attractive hypothesis is that persistent activity arises from a neural circuitry capable of temporal integration of its inputs, and might flexibly subserve a multitude of functions. Therefore, characterizing persistent activity in more detail under a diverse set of behavioral conditions and understanding its underlying biophysical mechanisms represent important

Box 2. Questions for future research

- What are the precise functions of persistent activity during working memory? Does the persistent activity observed in different brain areas perform different functions, such as selection and maintenance of remembered items? Does persistent activity reflect active control processes as well as maintenance, and how can we characterize these functions separately?
- How is persistent activity modified to store information about multiple sensory items that are presented simultaneously or sequentially [69]? How do the inputs to the neural circuits generate persistent activity gated so that it stores only behaviorally relevant information?
- What are the mechanisms of persistent activity? Does persistent activity depend on intracellular signals, short-term synaptic plasticity or feedback loops formed between multiple brain regions, including the basal ganglia and thalamus?
- What is the relationship between persistent activity and synaptic plasticity? Does persistent activity play a critical role in inducing synaptic plasticity and hence long-term memory? Conversely, what is the contribution of short-term synaptic plasticity to persistent activity?
- How is the time course of persistent activity controlled? Does persistent activity related to the animal's previous choices and their outcomes depend on the statistics of the task and environment? How is persistent activity terminated?
- What are the functions of neuromodulators, such as dopamine and norepinephrine, in controlling the strength and time course of persistent activity [70]?

goals if we are to uncover the biological basis of cognition (Box 2).

Acknowledgements

We are grateful to Hyojung Seo, Min Whan Jung and Xiao-Jing Wang for helpful discussions. This research was supported by National Institutes of Health Grants R01-EY016407 (C.E.C.), RL1-DA024855, R01-MH05916 and P05-NS048328 (D. L.).

References

- 1 Fuster, J.M. and Alexander, G.E. (1971) Neuron activity related to short-term memory. *Science* 173, 652–654
- 2 Kubota, K. and Niki, H. (1971) Prefrontal cortical unit activity and delayed alternation performance in monkeys. *J. Neurophysiol.* 34, 337–347
- 3 Funahashi, S. *et al.* (1989) Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J. Neurophysiol.* 61, 331–349
- 4 Baddeley, A. (1986) *Working Memory*, Oxford University Press
- 5 Goldman-Rakic, P.S. (1995) Cellular basis of working memory. *Neuron* 14, 477–485
- 6 Curtis, C.E. and D'Esposito, M. (2003) Persistent activity in the prefrontal cortex during working memory. *Trends Cogn. Sci.* 7, 415–423
- 7 Curtis, C.E. and D'Esposito, M. (2006) Working memory. In *Handbook of Functional Neuroimaging of Cognition* (2nd edn) (Cabezza, R. and Kingstone, A., eds), pp. 269–306, MIT Press
- 8 Schluppeck, D. *et al.* (2006) Sustained activity in topographic areas of human posterior parietal cortex during memory-guided saccades. *J. Neurosci.* 26, 5098–5108
- 9 Srimal, R. and Curtis, C.E. (2008) Persistent neural activity during the maintenance of spatial position in working memory. *Neuroimage* 39, 455–468
- 10 Curtis, C.E. *et al.* (2004) Maintenance of spatial and motor codes during oculomotor delayed response tasks. *J. Neurosci.* 24, 3944–3952
- 11 Todd, J.J. and Marois, R. (2004) Capacity limit of visual short-term memory in human posterior parietal cortex. *Nature* 428, 751–754
- 12 Vogel, E.K. *et al.* (2005) Neural measures reveal individual differences in controlling access to working memory. *Nature* 438, 500–503
- 13 Xu, Y.D. and Chun, M.M. (2006) Dissociable neural mechanisms supporting visual short-term memory for objects. *Nature* 440, 91–95
- 14 Baeg, E.H. *et al.* (2003) Dynamics of population code for working memory in the prefrontal cortex. *Neuron* 40, 177–188

- 15 Averbeck, B.B. and Lee, D. (2004) Coding and transmission of information by neural ensembles. *Trends Neurosci.* 27, 225–230
- 16 Averbeck, B.B. and Lee, D. (2007) Prefrontal neural correlates of memory for sequences. *J. Neurosci.* 27, 2204–2211
- 17 Hikosaka, O. *et al.* (2000) Role of the basal ganglia in the control of purposive saccadic eye movements. *Physiol. Rev.* 80, 953–978
- 18 Watanabe, Y. and Funahashi, S. (2004) Neuronal activity throughout the primate mediodorsal nucleus of the thalamus during oculomotor delayed-responses. I. Cue-, delay- and response-period activity. *J. Neurophysiol.* 92, 1738–1755
- 19 Wang, X.-J. (2008) Decision making in recurrent neuronal circuits. *Neuron* 60, 215–234
- 20 Seo, H. *et al.* (2007) Dynamic signals related to choices and outcomes in the dorsolateral prefrontal cortex. *Cereb. Cortex* 17, i110–i117
- 21 Chafee, M.V. and Goldman-Rakic, P.S. (1998) Matching patterns of activity in primate prefrontal area 8a and parietal area 7ip neurons during a spatial working memory task. *J. Neurophysiol.* 79, 2919–2940
- 22 Courtney, S.M. *et al.* (1998) An area specialized for spatial working memory in human frontal cortex. *Science* 279, 1347–1351
- 23 Miller, E.K. *et al.* (1991) A neural mechanism for working and recognition memory in inferior temporal cortex. *Science* 254, 1377–1379
- 24 Ó Scalaidhe, S.P. *et al.* (1999) Face-selective neurons during passive viewing and working memory performance of rhesus monkeys: evidence for intrinsic specialization of neuronal coding. *Cereb. Cortex* 9, 459–475
- 25 Ranganath, C. *et al.* (2004) Inferior temporal, prefrontal, and hippocampal contributions to visual working memory maintenance and associative memory retrieval. *J. Neurosci.* 24, 3917–3925
- 26 Jonides, J. *et al.* (1998) The role of parietal cortex in verbal working memory. *J. Neurosci.* 18, 5026–5034
- 27 Alain, C. *et al.* (2008) The contribution of the inferior parietal lobe to auditory spatial working memory. *J. Cogn. Neurosci.* 20, 285–295
- 28 Tark, K.J. and Curtis, C.E. (2009) Persistent neural activity in the human frontal cortex when maintaining space that is off the map. *Nat. Neurosci.* 12, 1463–1468
- 29 Romo, R. *et al.* (1999) Neuronal correlates of parametric working memory in the prefrontal cortex. *Nature* 399, 470–473
- 30 Fuster, J.M. (2001) The prefrontal cortex—an update: time is of the essence. *Neuron* 30, 319–333
- 31 Passingham, D. and Sakai, K. (2004) The prefrontal cortex and working memory: physiology and brain imaging. *Curr. Opin. Neurobiol.* 14, 163–168
- 32 Wise, S.P. (2008) Forward frontal fields: phylogeny and fundamental function. *Trends Neurosci.* 31, 599–608
- 33 Meyer, T. *et al.* (2007) Persistent discharges in the prefrontal cortex of monkeys naive to working memory tasks. *Cereb. Cortex* 17, i70–i76
- 34 Andersen, R.A. and Buneo, C.A. (2002) Intentional maps in posterior parietal cortex. *Annu. Rev. Neurosci.* 25, 189–220
- 35 Curtis, C.E. and Connolly, J.D. (2008) Saccade preparation signals in the human frontal and parietal cortices. *J. Neurophysiol.* 99, 133–145
- 36 Rainer, G. *et al.* (1999) Prospective coding for objects in primate prefrontal cortex. *J. Neurosci.* 19, 5493–5505
- 37 Freedman, D.J. *et al.* (2001) Categorical representation of visual stimuli in the primate prefrontal cortex. *Science* 291, 312–316
- 38 Wallis, J.D. *et al.* (2001) Single neurons in prefrontal cortex encode abstract rules. *Nature* 411, 953–956
- 39 Bunge, S.A. *et al.* (2003) Neural circuits subserving the retrieval and maintenance of abstract rules. *J. Neurophysiol.* 90, 3419–3428
- 40 Hempel, C.M. *et al.* (2000) Multiple forms of short-term plasticity at excitatory synapses in rat medial prefrontal cortex. *J. Neurophysiol.* 83, 3031–3041
- 41 Mongillo, G. *et al.* (2008) Synaptic theory of working memory. *Science* 319, 1543–1546
- 42 Gold, J.I. and Shadlen, M.N. (2007) The neural basis of decision making. *Annu. Rev. Neurosci.* 30, 535–574
- 43 Seo, H. and Lee, D. (2007) Temporal filtering of reward signals in the dorsal anterior cingulate cortex during a mixed-strategy game. *J. Neurosci.* 27, 8366–8377
- 44 Laming, D.R.J. (1968) *Information Theory of Choice–Reaction Times*, Academic Press
- 45 Roitman, J.D. and Shadlen, M.N. (2002) Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.* 22, 9475–9489
- 46 Heekeren, H.R. *et al.* (2008) The neural systems that mediate human perceptual decision making. *Nat. Rev. Neurosci.* 9, 467–479
- 47 Barraclough, D.J. *et al.* (2004) Prefrontal cortex and decision making in a mixed-strategy game. *Nat. Neurosci.* 7, 404–410
- 48 Seo, H. *et al.* (2009) Lateral intraparietal cortex and reinforcement learning during a mixed-strategy game. *J. Neurosci.* 29, 7278–7289
- 49 Seo, H. and Lee, D. (2009) Behavioral and neural changes after gains and losses of conditioned reinforcers. *J. Neurosci.* 29, 3627–3641
- 50 Histed, M.H. *et al.* (2009) Learning substrates in the primate prefrontal cortex and striatum: sustained activity related to successful actions. *Neuron* 63, 244–253
- 51 Kim, H. *et al.* (2009) Role of striatum in updating values of chosen actions. *J. Neurosci.* 29, 14701–14712
- 52 Sutton, R.S. and Barto, A.G. (1998) *Reinforcement Learning: An Introduction*, MIT Press
- 53 Lee, D. *et al.* (2004) Reinforcement learning and decision making in monkeys during a competitive game. *Cogn. Brain Res.* 22, 45–58
- 54 Samejima, K. *et al.* (2005) Representation of action-specific reward values in the striatum. *Science* 310, 1337–1340
- 55 Behrens, T.E.J. *et al.* (2007) Learning the value of information in an uncertain world. *Nat. Neurosci.* 10, 1214–1221
- 56 Schultz, W. (2006) Behavioral theories and the neurophysiology of reward. *Annu. Rev. Psychol.* 57, 87–115
- 57 Lee, D. (2006) Neural basis of quasi-rational decision making. *Curr. Opin. Neurobiol.* 16, 191–198
- 58 Daw, N.D. and Doya, K. (2006) The computational neurobiology of learning and reward. *Curr. Opin. Neurobiol.* 16, 199–204
- 59 Kim, S. *et al.* (2008) Prefrontal coding of temporally discounted values during inter-temporal choice. *Neuron* 59, 161–172
- 60 Lau, B. and Glimcher, P.W. (2008) Value representations in the primate striatum during matching behavior. *Neuron* 58, 451–463
- 61 Sugrue, L.P. *et al.* (2004) Matching behavior and the representation of value in the parietal cortex. *Science* 304, 1782–1787
- 62 Walton, M.E. *et al.* (2010) Separable learning systems in the macaque brain and the role of orbitofrontal cortex in contingent learning. *Neuron* 65, 927–939
- 63 Pan, W.-X. *et al.* (2005) Dopamine cells respond to predicted events during classical conditioning: evidence for eligibility traces in the reward-learning network. *J. Neurosci.* 25, 6235–6242
- 64 Kim, Y. *et al.* (2007) Encoding of action history in the rat ventral striatum. *J. Neurophysiol.* 98, 3548–3556
- 65 Ito, M. and Doya, K. (2009) Validation of decision-making models and analysis of decision variables in the rat basal ganglia. *J. Neurosci.* 29, 9861–9874
- 66 Matsumoto, M. *et al.* (2007) Medial prefrontal cell activity signaling prediction errors of action values. *Nat. Neurosci.* 10, 647–656
- 67 Fransén, E. *et al.* (2006) Mechanism of graded persistent cellular activity of entorhinal cortex layer V neurons. *Neuron* 49, 735–746
- 68 Zucker, R.S. and Regehr, W.G. (2002) Short-term synaptic plasticity. *Annu. Rev. Physiol.* 64, 355–405
- 69 Warden, M.R. and Miller, E.K. (2007) The representation of multiple objects in prefrontal neuronal delay activity. *Cereb. Cortex* 17, i41–i50
- 70 Robbins, T.W. and Arnsten, A.F.T. (2009) The neuropsychopharmacology of fronto-executive function: monoaminergic modulation. *Annu. Rev. Neurosci.* 32, 267–287