

# Basic mechanisms for graded persistent activity: discrete attractors, continuous attractors, and dynamic representations

Carlos D Brody<sup>\*†</sup>, Ranulfo Romo<sup>‡</sup> and Adam Kepecs<sup>\*</sup>

Persistent neural activity is observed in many systems, and is thought to be a neural substrate for holding memories over time delays of a few seconds. Recent work has addressed two issues. First, how can networks of neurons robustly hold such an active memory? Computer systems obtain significant robustness to noise by approximating analogue quantities with discrete digital representations. In a similar manner, theoretical models of persistent activity in spiking neurons have shown that the most robust and stable way to store the short-term memory of a continuous parameter is to approximate it with a discrete representation. This general idea applies very broadly to mechanisms that range from biochemical networks to single cells and to large circuits of neurons. Second, why is it commonly observed that persistent activity in the cortex can be strongly time-varying? This observation is almost ubiquitous, and therefore must be taken into account in our models and our understanding of how short-term memories are held in the cortex.

## Addresses

<sup>\*</sup>Cold Spring Harbor Laboratory, 1 Bungtown Road, Cold Spring Harbor, NY 11724, USA

<sup>†</sup>e-mail: brody@cshl.edu

<sup>‡</sup>Instituto de Fisiología Celular, Universidad Nacional Autónoma de México, 04510, México DF, México

**Current Opinion in Neurobiology** 2003, **13**:204–211

This review comes from a themed issue on  
Cognitive neuroscience  
Edited by Brian Wandell and Anthony Movshon

0959-4388/03/\$ – see front matter  
© 2003 Elsevier Science Ltd. All rights reserved.

DOI 10.1016/S0959-4388(03)00050-3

## Abbreviations

**L** Lyapunov  
**PFC** prefrontal cortex

## Introduction

The problem at the core of memory formation is to establish a representation of a stimulus that persists after the stimulus is gone. When neural activity is not only stimulus-selective but also persistent, outlasting the stimulus that initially drove it, it can be used as a memory representation of the stimulus. Such persistent, stimulus-dependent activity is thought to be the neural substrate for memories that last anywhere up to tens of seconds [1–3], and has been observed in a wide variety of regions of the nervous system [4–6]. For example, in the

prefrontal cortex of monkeys that are trained to perform short-term memory tasks, persistent neural activity has been interpreted as the basis of working memory, that is, the ability to hold something ‘in mind’ for a few seconds [1,3,7–10]. In the anterior dorsal nucleus of the rodent thalamus (as well as in other brain areas), persistent activity is thought to allow encoding of a memory of head direction [11,12]. Additionally, in area I of the goldfish hindbrain, which controls the motor system that drives the eye muscles, persistent activity is interpreted as a short-term memory of eye position that keeps the eyes still between saccades [13–15].

For persistent neural activity to serve as memory, it must be robust to distractors and noise, at least over the period during which the organism requires the memory. Robustness to noise is thus a key issue for memory. In a biological context, the memory system that generates persistent activity must also be robust to significant variability in the components used to build it (even when the components are noise-free). For example, it would be unrealistic to expect a memory network to require 1% precision in synaptic connection strengths. Another key issue for biological memory, then, is robustness to imprecise components.

The two issues of robustness to noise and robustness to imprecise components are important to all biological memory systems at all levels from the subcellular to the cellular to the cortical network. The first part of this review presents the general conceptual framework that underlies current and upcoming research in this area. This framework, which uses concepts borrowed from dynamical systems theory, is important because it applies very broadly to all systems dedicated to memory.

When it comes to the particular case of persistent cortical activity, however, the framework faces important challenges. Persistent activity is often found in cortical areas that are involved in more than just the pure maintenance of information. For example, persistent activity co-exists with decision making and motor planning in the lateral intraparietal cortex [4,16,17<sup>\*</sup>], the frontal eye fields [18] and the premotor cortex [19]. Persistent activity also co-exists with expected reward coding and anticipation in the prefrontal cortex [20–22]. In contrast to the static representation assumed in the framework presented in the first part of this review, the representation of short-term memory information in prefrontal and other cortices is often found to be strongly dynamic. That is to say that firing rates, although persistently stimulus-dependent,

can change markedly and systematically during the memory period. The second part of this review underscores the observation of these dynamics: their ubiquity suggests that rather than being peripheral to the memory process, they may be central to it. These systematic dynamics should therefore be incorporated into cortical models of short-term memory.

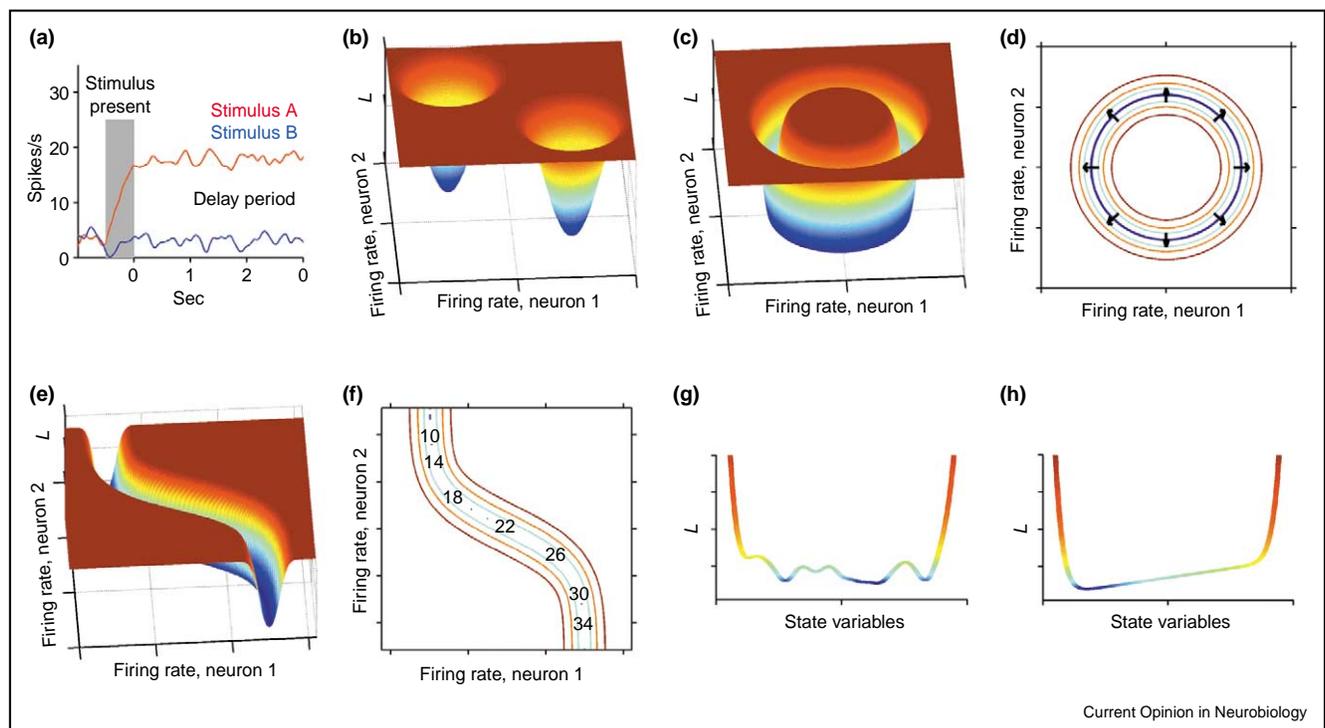
### Robustness to noise and imprecise components

Stimulus-dependence in persistent activity is thought to correspond to a state of the neural memory system in which the system will remain (persist) even after the stimulus is gone (Figure 1a). This state must be stable, in the sense that the system must remain in this state even if disturbed slightly by noise. For this to occur, there must be forces that restore the system back to the stable state if noise makes it stray slightly away. Let us imagine that there is a function of the state of the system, which we shall call  $L$ , and that the dynamics of the system are such that  $L$  always tends to fall in value as time progresses. The restoring forces around the stable point can then be visualized as walls of  $L$  that form a well around that point. Because  $L$  tends to decrease, once the system is in a state inside the well, it will tend to decline to the lowest level

of the well and then stay there (Figure 1b). At least locally around any stable fixed point,  $L$  can always be defined. Hence we can always use it for the purpose of visualization, even if further away from the stable points (e.g. in the reddish brown regions in Figure 1) the precise value of  $L$  may not actually be defined. A network of neurons with two different persistent stable states can thus be visualized as an  $L$  function (technically known as a *Lyapunov* function) with two separate wells (Figure 1b). By identifying the position the system's state is in, it is possible to decode the memory held in the system: that is, if in well A, the memory is 'A'; if in well B, the memory is 'B'.

A system with discrete wells can store the memory of one of a set of discrete items [23–25]. However, parameters that can take on a continuous range of values can also be stored in short-term memory. For example, monkeys can be trained to remember the intended direction of a saccade when the response is delayed for a few seconds [4,7,26]. This direction can, in principle, range from  $0^\circ$  to  $360^\circ$ . To visualize this as an  $L$ -function diagram, imagine that the monkey is holding '45°' in memory. There must therefore be a well for  $45^\circ$ . As a position infinitesimally close to  $45^\circ$  is also a possible memory, there must be

Figure 1



Attractor shapes for persistent activity. (a) A schematic of the firing rate of a persistent, stimulus-dependent neuron. (b) An  $L$  function with two wells. Here, the state variables ( $x$  and  $y$  axes) are plotted as the firing rates of two neurons but, in general, they are all the variables necessary to define the system's state. They could thus include, for example,  $\text{Ca}^{2+}$  concentration in dendrites. (c) A ring attractor (see text). (d) Contour plot for the ring attractor (c). Each position in the ring codes for a different direction. (e) A line attractor. (f) Contour plot for the line attractor (e). Each position on the line codes for a value of the stimulus being held in memory (e.g. frequency of a vibrotactile stimulus [10]). (g) and (h) Slight imperfections (g) or a slight tilt (h) destroy the flat bottom of a continuous attractor.

another well, infinitesimally close to the  $45^\circ$  one; and there must be another well infinitesimally close to that one. This must continue all the way around the circle of directions, to  $360^\circ$ , and then back round to  $45^\circ$ . In short, there must be a continuum of wells (i.e. minima) of  $L$  that loops back onto itself. The only way for this to occur is for the continuous set of wells to carve out a valley that forms a loop (Figure 1c). Different positions along the bottom of this valley, called a ‘ring attractor’, represent all the different possible memories for direction [27–32].

Not all continuous parameters are defined in a periodic space (that is, a space that loops back onto itself, in the sense that  $360^\circ$  is the same as  $0^\circ$ ). For example, subjects may be asked to remember the frequency of a vibrotactile stimulus [10]. Vibration frequency is continuous but it is not defined in a periodic space. In this case, the continuum of minima in the  $L$ -function diagram forms a valley that doesn’t loop back on itself (Figure 1e). Valleys of this kind are called ‘line attractors’. This conceptual approach, as applied to the short-term memory of eye position, was clearly described by Seung [14].

### Consequences of perfect continuous attractors

The  $L$ -function diagram allows us to deduce important general properties of continuous attractors. First, as described above, all points along the line that describes the bottom of the attractor must be minima of  $L$ . This means that all of these points must all have the same value of  $L$ , as no point could be lower than any of its neighbors. In turn, this leads us to the conclusion that there are no restorative forces along the bottom of the valley — the walls of  $L$  represent the restorative forces. Random noise is thus equally likely to move the system up or down the line of a continuous attractor, and the system moves in a random walk (diffusion). Systems that are based on networks of large numbers of neurons may reduce noise by averaging over neurons, and therefore reduce the speed with which the system state randomly drifts away from its original position. Despite this, the speed of the random walk will never be exactly zero.

The construction of continuous attractors presents significant difficulties. The requirement that  $L$  be exactly the same along the bottom of the valley is, in essence, a perfect symmetry requirement: circular symmetry in the case of the ring attractor (Figure 1c), and translational symmetry in the case of the line attractor (Figure 1e). Any departure from this perfect symmetry can substantially perturb a continuous attractor. For example, slight imperfections along the bottom of the attractor (Figure 1g) will mean that only a few points along the line are minima of  $L$ , and therefore stable. Even a very slight tilt on the overall shape of the attractor would lead to a steady drift down its slope (Figure 1h). Different plasticity mechanisms may place the onus for symmetry on different

parameters, but it is an onus that cannot be entirely avoided. Recent theoretical work has suggested that it may be advantageous to place the symmetry onus on cellular homeostasis parameters [33].

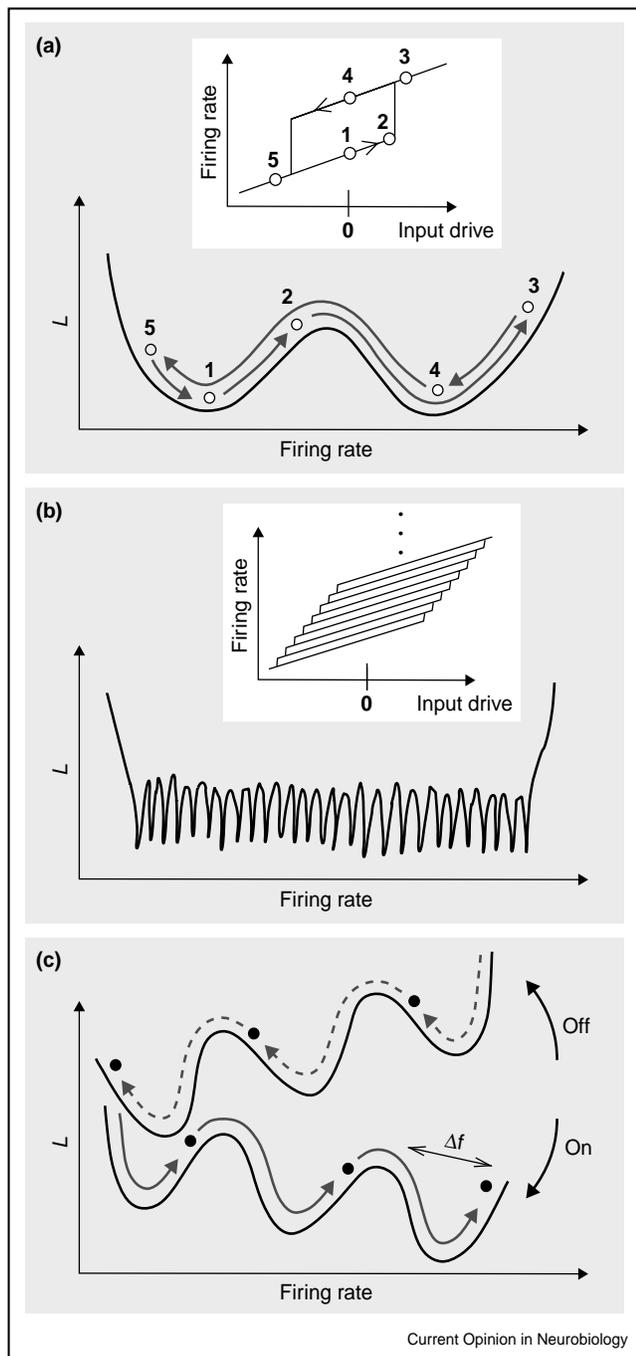
Determining how non-linear spiking neurons could interconnect so as to form a line attractor is a difficult problem. In a landmark study [34], Seung and co-workers devised a method to design a spiking neural network model that approximates a line attractor. Their model successfully replicated the main features of data recorded from persistent neurons that encode eye position in goldfish. However, this network also highlighted the construction problems involved in building a line attractor. Changing connection strength between the neurons by as little as 1% produced unacceptably rapid drifts in the predicted eye position. In other words, even a slight tilt in the attractor (Figure 1h) led the model to predict results that were incompatible with behavior in the fish. Such precision in construction appears to be biologically unrealistic.

Recently, Koulakov *et al.* [35\*\*] proposed a new class of models that are based on many individually hysteretic units. They combined analogue and digital aspects in an attempt to alleviate both problems of robustness in construction and robustness to noise. In their model, a modification of up to 16% in connection strengths can be allowed without loss of function. The ideas underlying the robustness of this model are quite general, and form the basis of several other models that are currently being developed.

### Relationship between hysteresis and multistability

Hysteretic and multistable systems are intimately related. The simplest multistable system that we can imagine is a bistable one. In Figure 2a we show an example of an  $L$ -function diagram with two wells. (For clarity of exposition, we have plotted  $L$  as a function of only one state variable, which could, for example, be the firing rate of a neuron.) Suppose the system is at rest, in the state marked ‘1’. Now imagine that as an input (e.g. an injected current) increases, the system gets pulled to the right in the  $L$ -function diagram, perhaps far enough to go all the way up to the lip of the well (state 2). As shown in the inset, firing rate increases smoothly from state 1 to state 2 as the input drive increases. However, if the input drive is strong enough to push the system over the lip of the well, the system will undergo an abrupt transition, moving from the right edge of the first well (state 2) to the right edge of the second well (state 3). There would, thus, be an abrupt change of state accompanied by an abrupt increase in firing rate. After this transition, releasing the input drive will not bring the system back to state 1, but will leave it at the bottom of the second well (state 4). To bring the system back to state 1, a significant negative input drive is required, a negative input drive that is enough to pull the

Figure 2



Hysteresis, bistability, and the discrete approximation to a continuous attractor. **(a)** A system that has two wells in  $L$  will show hysteresis (inset). **(b)** Many individual wells, arranged along a line, are associated with many hysteresis loops (inset). The location of an arbitrary position is approximated by choosing the nearest well, and using that to encode the position. These wells are robust to disruptions such as the noise or tilt illustrated in Figure 1g,h. **(c)** A force to the right is visualized as a tilt of  $L$  to the right. As it moves along the well, the system spends more time on the right-hand side than on the left. This is conversely true for a force in the opposite direction.

system back over the barrier that separates the two wells. When the firing rate is plotted as a function of input drive, we find that the firing rate depends on the history of the system, creating the two curves in the inset of Figure 1b, one for increasing input drive, and the other for decreasing input drive. Behavior that is history-dependent, and remains so over long timescales, is the defining feature of hysteretic systems.

### The hysteretic continuum

A bistable system can only store a binary variable (the choice of one of two wells), not a continuous one. Compared to a continuous attractor, however, a bistable system is very robust. The same slight disturbances that can destroy the continuity of an attractor valley (represented in Figure 1g,h as adding a little noise or tilt to the  $L$  function) do not destroy the fundamental properties of the bistable system, that is, the existence of two separate wells. Can we keep the robustness properties of the bistable system while allowing the storage of a continuous variable? The answer to this question is 'no'; as we have shown in the previous section, continuity imposes properties of symmetry on  $L$  that necessarily render the system fragile to disturbances. It is possible, however, to robustly store a discrete approximation to a continuous variable.

Imagine a system in which there are many wells in the  $L$  function rather than just two (Figure 2b). Suppose that these wells are densely distributed over the same range of firing rates as the two wells of Figure 2a. Instead of there being a single hysteretic loop, as in the inset of Figure 2a, there are now many hysteretic loops, each of which describes a small jump in the firing rate (inset of Figure 2b). The deeper the wells are, the wider the hysteretic loop, and the more robust the system is to disturbances. Even if the wells are deep, however, they may be spaced very close to each other. Thus, although any finite system will have a finite number of wells, the wells may be so closely spaced as to allow the storage of a very good approximation of a continuous variable. The most familiar example of a non-neural hysteretic system that appears to be continuous is that of digital computers. We routinely use digital computers to represent continuous variables — for example, when simulating Hodgkin–Huxley models — yet internally, the representation of continuous variables is composed of a very high density of discrete bits, each of which is stored as the state of a bistable microsystem (a flip/flop).

### Generality of the hysteretic continuum concept

There are many different neural mechanisms that could be used to implement a system of closely spaced wells, as depicted in the schematic of Figure 2b. Theoretical models have proposed forming a single large system with many stable states out of a large collection of individually bistable units that are weakly coupled together. (One of

the challenges, which we do not explore here, is to structure the coupling between subunits so that the many wells of the system are arranged into a single line; it is only then that the system can be used as an integrator.) Koulakov *et al.* [35\*\*] initially proposed two possible models for the formation of bistable units at the scale of single cells or larger. One possibility used strong coupling within small groups of neurons to make each group bistable. The other possibility used the voltage-dependence of NMDA channels to make single neurons bistable [36].

Bistable subunits may also be formed at a scale smaller than a single neuron. Recent *in vitro* work by Egorov and co-workers [37\*\*] has shown that individual neurons in entorhinal cortex slices are capable of displaying multistable graded persistent activity, as long as the right chemical environment is provided in the bath. This remarkable result has been followed by theoretical work in which individual bistable subunits are formed from compartments of dendritic branches that contain voltage-dependent  $\text{Ca}^{2+}$  channels [38,39]. Using a large number of such bistable compartments per neuron, models of individual neurons may be constructed to display many different stable states, each with a different stable persistent firing rate. One proposal arranged all the compartments in a single line, as part of a single dendritic shaft. This spatial arrangement was then part of the constraints that were used to create a single line of densely arranged wells [39].

The power of the  $L$  function approach lies in its generality: the basic concepts and results of Figures 1 and 2 are, in essence, geometric. Thus, the idea that the storage of a continuous variable is more robustly achieved using a discrete approximation is one that applies across many spatial and temporal scales, and across a broad range of systems. In particular, this idea applies to memory systems that range from the biochemical basis of LTP in single spines [40,41] to seconds of persistent activity in single neurons [37\*\*,39,42], to biochemical cellular memories in non-neural cells [43], to networks of neurons [35\*\*], and to modern computer memories. Hysteresis and its link to robustness are likely to be central concepts in research into persistent activity that codes for the memory of continuous variables.

If the  $L$ -function concept is so general that it is not specific to a particular mechanism, how can it be tested? One general prediction is based on the existence of wells with a finite (non-zero) width. Let us visualize input pushing the system in the 'ON' or 'OFF' directions by tilting the  $L$  function, either to the right ('ON') or to the left ('OFF'). If the system is moving smoothly in the 'ON' direction, it will tend to spend most of its time on the right-hand side of the wells; if moving smoothly in the 'OFF' direction, it will spend most of its time on the left-hand side of the wells. This will lead to a systematic difference in firing rates ( $\Delta f$ ) when the system is moving

in one direction compared to the other (Figure 2c): that is, to a system-wide hysteresis [35\*\*].

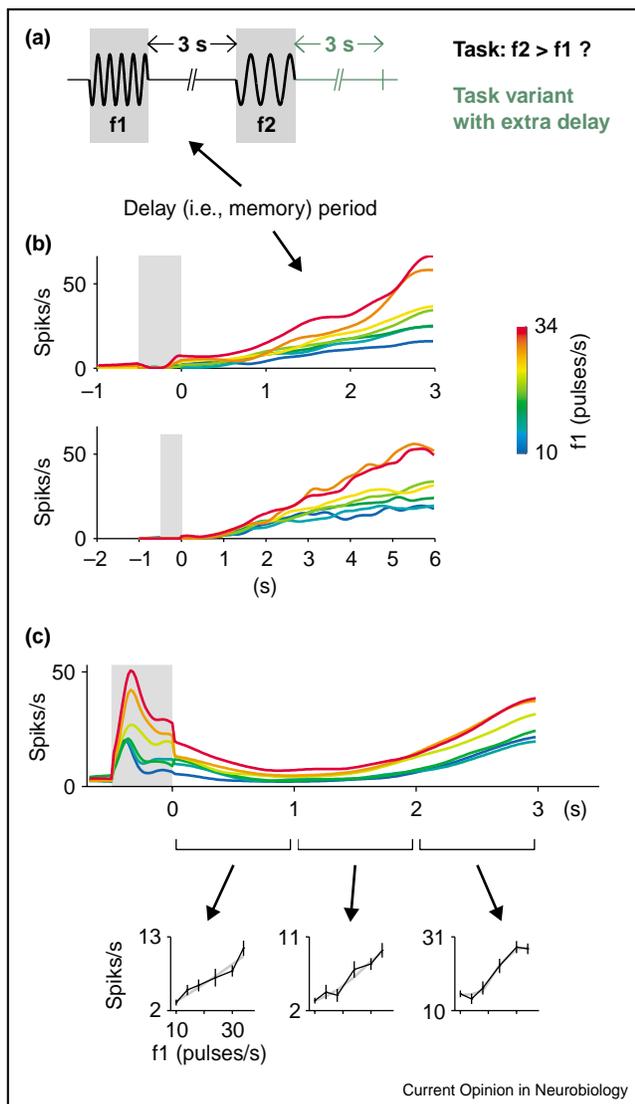
### Dynamic representations of short-term memory

We now turn to the second subject of this review. So far, we have considered conceptual approaches to understanding stimulus-dependent activity that is 'statically' persistent. Once such a system is in a state that produces a particular set of firing rates, the system stays there and firing rates are fixed until an external signal terminates the delay period. Assigning a meaning to a static and stable system state is an attractive and simple approach to conceptualizing the memory problem.

Nevertheless, we have known since the first reports of persistent neural activity in the cortex [44,45] that the firing rates of persistently active cortical neurons often change systematically over the course of delay periods. For example, Brody *et al.* [46\*] trained monkeys to remember a continuous scalar parameter — the frequency of a mechanical vibration applied to a fingertip — over delay periods that lasted a few seconds (Figure 3a). They then recorded the responses of neurons in the prefrontal cortex (PFC). Most of the neurons that fired during the stimulus-dependent delay period had a strongly time-varying response profile. Three-quarters of the neurons were stimulus-dependent only during a part of the delay period, generally either towards the end (33% of neurons, dubbed 'late' neurons; Figure 3b) or the beginning (31%, 'early' neurons, not shown). Only one-quarter of the neurons had firing rates that persisted in being stimulus-dependent throughout the entire delay period ('persistent' neurons; Figure 3c). This pattern of results might lead one to imagine that during long delay periods, there are stretches of time, far from the start of the delay period and far from its end, during which only the 'persistent' class of neurons have stimulus-dependent activity. Thus, the apparently more static 'persistent' neurons might be a special set of neurons that are uniquely responsible for actively maintaining a memory during long delay periods. However, when 'late' neurons were recorded both during a block of trials with a 3-s delay period and during a subsequent block of trials with a 6-s delay, they were found to stretch the timing of their responses [46\*]. This meant that the evolution of their firing profiles took twice as long in trials with a 6-s delay period (Figure 3b, lower panel) than in trials involving a 3-s delay period (Figure 3b, upper panel). Thus, long delay periods do not isolate moments in time during which only the 'persistent' neurons carry stimulus-dependent information. This suggests that all three types of neurons, 'early', 'persistent,' and 'late', could play equally important roles in representing and maintaining a memory.

The firing rates of 'persistent' neurons vary markedly and systematically with time (e.g. Figure 3c [note  $y$  axes in

Figure 3



Neurons in the prefrontal cortex (PFC) with persistent activity that is both stimulus-dependent and systematically time-varying [10,46\*]. (a) The monkeys' task was to compare the frequencies of two mechanical vibrations applied sequentially to a fingertip, and decide which was the higher. In the standard task, the monkeys could respond immediately after  $f_2$ , and then get their reward if they were correct. In a variant of the task, they were required to wait for an additional 3 s before responding. (b) (Upper panel) Firing rate of a 'late' neuron during the  $f_1$  stimulus and delay period. The color code on the right indicates the value of the  $f_1$  stimulus. (Lower panel) The same neuron was recorded during a subsequent block of trials that had a 6-s delay period between  $f_1$  and  $f_2$ . (c) A 'persistent' neuron.

lower panels]). A remarkably high fraction of the 'persistent' neurons recorded by Brody *et al.* [46\*], about 90%, had firing rates that varied significantly over time during the delay period ( $p < 0.001$  in an ANOVA test). Only 3% of neurons had reliably static ( $p > 0.1$ ) firing rates. In addition, neurons that had time-varying firing rates

(whether members of the 'persistent' or 'late' classes) were as predictive of the animal's ultimate response choice as the few neurons with approximately static persistent activity [46\*]. Time-dependent neurons thus seem to play a similar role to that played by time-independent neurons in determining memory-based behavior. Many other researchers have also reported time-varying firing of neurons during delay periods (e.g. [9,16,44,47\*]). This pattern of firing appears to be an almost ubiquitous characteristic of stimulus-dependent delay period activity in the cortex. Taken together, these results suggest that, in the cortex, the short-term memory of the stimulus is mostly represented in neurons with dynamic, not static, firing rates.

If the neurons that represent the memory of the stimulus have firing rates that change with time, does this imply that the memory itself is also changing? Not necessarily. For example, the sum of the firing rates of neurons with opposing time-dependent properties could produce a static stimulus-dependent representation. Additionally, the time-dependence of the firing rates could perhaps explicitly encode time itself [48]. Knowing how much time had elapsed since the stimulus would allow us to determine the appropriate stimulus→firing-rate map (i.e. the choice of lower panel in Figure 3c), and would therefore allow us to correctly decode the stimulus from the firing rate.

### Continuous stimulus recoding

Some neurons in the PFC are known to code for an expected reward, as well as for the memory of a past stimulus [20,21,22,49]. This activity is often of a 'ramping-up' nature. When ramping-up, anticipatory encoding is multiplexed with stimulus-dependent encoding, which could lead to responses such as those illustrated in Figure 3b. However, in a variant of their standard frequency discrimination task, Brody *et al.* [46\*] placed an extra 3-s delay between the stimulus  $f_2$  and the reward (green in Figure 3a). In this situation, there was no immediate reward or motor act at the end of the first delay period, and no trial-by-trial variation in the reward to be encoded. Nevertheless, neurons with the same temporal profiles of firing rate and stimulus-dependence were found in the first delay period of the variant task and in the delay period of the standard task. This included ramping-up neurons and neurons with 'late' stimulus dependence. Such neurons are thus not necessarily coding for the anticipation of an immediately upcoming reward.

It has been suggested that the encoding of short-term memories (separately from reward coding) in the PFC may shift during the delay period from an encoding of the stimulus itself (retrospective coding) to an encoding of a stimulus-dependent expectation of what will occur at the end of the delay period (prospective coding) [50,51\*]. These two types of encoding, retrospective

and prospective, can be associated with ramping down and ramping up activity, respectively. They may be shared in the same neuron in different proportions at different times during the delay period. This coding shift has been termed ‘cross-temporal integration’ [51\*].

A general description that could apply across several of these different types of time-dependent responses is that, during the delay period, the encoding of the stimulus may be actively transformed into a representation that facilitates the task requirements of the end of the delay period, whatever these may be. These requirements range from reward coding for evaluation of reward and behavior monitoring in some neurons and tasks to stimulus recoding for comparison with upcoming stimuli (e.g. comparison of stimuli f1 and f2 in Figure 3c) in others. The key point about this view is that the goal of cortical short-term memory processes is not seen as the static storage of information about previous stimuli (as in the first part of this review). Instead, it is seen as an essentially dynamic process that uses stimulus information to anticipate and, most importantly, prepare the organism for upcoming behavioral and cognitive task requirements.

### Conclusions and outlook

Imparting robustness to graded memory systems by building them out of many individually bistable hysteretic subunits is a promising idea for understanding the basic mechanisms behind the active storage of short-term memories of continuous variables.

In the cortex, however, persistent activity is often characterized by being not only stimulus-dependent but also time-varying. It may be that a few neurons with statically persistent stimulus-dependent firing rates (in cortex or elsewhere) are responsible for maintaining the representation of short-term memories. However, we suggest that it is more likely that the representation is fundamentally a dynamic one, with time-dependent neurons playing as central a part as time-independent (static) ones. Time-dependent firing may reflect participation in not just memory storage but also preparation for anticipated computations. Characterizing these computations, understanding their interaction with memory storage, and modeling their dynamics [47\*] is likely to be a necessary and important part of elucidating the cortical mechanisms for short-term memory [51\*].

### Acknowledgements

This work was supported by startup funds of Cold Spring Harbor Laboratory.

### References and recommended reading

Papers of particular interest, published within the annual period of review, have been highlighted as:

- of special interest
- of outstanding interest

1. Fuster JM: *Memory in the Cerebral Cortex*. Cambridge: MIT Press;1995.
2. Goldmanrakis PS: **Cellular basis of working-memory**. *Neuron* 1995, **14**:477-485.
3. Goldmanrakis PS: **Regional and cellular fractionation of working memory**. *Proc Natl Acad Sci USA* 1996, **93**:13473-13480.
4. Gnadt JW, Andersen RA: **Memory related motor planning activity in posterior parietal cortex of macaque**. *Exp Brain Res* 1988, **70**:216-220.
5. McFarland JL, Fuchs AF: **Discharge patterns in nucleus-prepositus-hypoglossi and adjacent medial vestibular nucleus during horizontal eye-movement in behaving macaques**. *J Neurophysiol* 1992, **68**:319-332.
6. Prut Y, Fetz EE: **Primate spinal interneurons show pre-movement instructed delay activity**. *Nature* 1999, **401**:590-594.
7. Funahashi S, Bruce CJ, Goldmanrakis PS: **Mnemonic coding of visual space in the monkeys dorsolateral prefrontal cortex**. *J Neurophysiol* 1989, **61**:331-349.
8. Miller EK, Li L, Desimone R: **Activity of neurons in anterior inferior temporal cortex during a short-term-memory task**. *J Neurosci* 1993, **13**:1460-1478.
9. Miller EK, Erickson CA, Desimone R: **Neural mechanisms of visual working memory in prefrontal cortex of the macaque**. *J Neurosci* 1996, **16**:5154-5167.
10. Romo R, Hernandez A, Lemus L, Zainos A, Brody CD: **Neuronal correlates of decision-making in secondary somatosensory cortex**. *Nat Neurosci* 2002, **5**:1217-1225.
11. Taube JS: **Head direction cells and the neurophysiological basis for a sense of direction**. *Prog Neurobiol* 1998, **55**:225-256.
12. Sharp PE, Blair HT, Cho JW: **The anatomical and computational basis of the rat head-direction cell signal**. *Trends Neurosci* 2001, **24**:289-294.
13. Pastor AM, Delacruz RR, Baker R: **Eye position and eye velocity integrators reside in separate brain-stem nuclei**. *Proc Natl Acad Sci USA* 1994, **91**:807-811.
14. Seung HS: **How the brain keeps the eyes still**. *Proc Natl Acad Sci USA* 1996, **93**:13339-13344.
15. Aksay E, Baker R, Seung HS, Tank DW: **Anatomy and discharge properties of pre-motor neurons in the goldfish medulla that have eye-position signals during fixations**. *J Neurophysiol* 2000, **84**:1035-1049.
16. Chafee MV, Goldman-Rakic PS: **Matching patterns of activity in primate prefrontal area 8a and parietal area 7ip neurons during a spatial working memory task**. *J Neurophysiol* 1998, **79**:2919-2940.
17. 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**:805-811.
 An interesting paper that describes experimental data showing that temporal structured activity (specifically, in the 25–90 Hz band of the local field potential) carries information about actively held short-term memories. This is different to the standard view (adopted in this review), which only considers the slow changes in firing rates, not those in fast  $\gamma$ -band temporal structure. The precise relation between the two is still unclear.
18. Kim JN, Shadlen MN: **Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque**. *Nat Neurosci* 1999, **2**:176-185.
19. Hernandez A, Zainos A, Romo R: **Temporal evolution of a decision-making process in medial premotor cortex**. *Neuron* 2002, **33**:959-972.
20. Watanabe M: **Reward expectancy in primate prefrontal neurons**. *Nature* 1996, **382**:629-632.
21. Hikosaka K, Watanabe M: **Delay activity of orbital and lateral prefrontal neurons of the monkey varying with different rewards**. *Cereb Cortex* 2000, **10**:263-271.
22. Kobayashi S, Lauwereyns J, Koizumi M, Sakagami M, Hikosaka O: **Influence of reward expectation on visuospatial processing in**

- macaque lateral prefrontal cortex. *J Neurophysiol* 2002, **87**:1488-1498.**
23. Hopfield JJ: **Neurons with graded response have collective computational properties like those of 2-state neurons.** *Proc Natl Acad Sci USA* 1984, **81**:3088-3092.
  24. Zipser D, Kehoe B, Littlewort G, Fuster J: **A spiking network model of short-term active memory.** *J Neurosci* 1993, **13**:3406-3420.
  25. Amit DJ, Brunel N: **Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex.** *Cereb Cortex* 1997, **7**:237-252.
  26. Funahashi S, Bruce CJ, Goldman-Rakic PS: **Visuospatial coding in primate prefrontal neurons revealed by oculomotor paradigms.** *J Neurophysiol* 1990, **63**:814-831.
  27. Benyishai R, Baror RL, Sompolinsky H: **Theory of orientation tuning in visual-cortex.** *Proc Natl Acad Sci USA* 1995, **92**:3844-3848.
  28. Zhang K: **Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory.** *J Neurosci* 1996, **16**:2112-2126.
  29. Camperi M, Wang XJ: **A model of visuospatial working memory in prefrontal cortex: recurrent network and cellular bistability.** *J Comput Neurosci* 1998, **5**:383-405.
  30. Compte A, Brunel N, Goldman-Rakic PS, Wang XJ: **Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model.** *Cereb Cortex* 2000, **10**:910-923.
  31. Gutkin BS, Laing CR, Colby CL, Chow CC, Ermentrout GB: **Turning on and off with excitation: the role of spike-timing asynchrony and synchrony in sustained neural activity.** *J Comput Neurosci* 2001, **11**:121-134.
  32. Laing CR, Chow CC: **Stationary bumps in networks of spiking neurons.** *Neural Comput* 2001, **13**:1473-1494.
  33. Renart A, Song S, Wang X-J: **Robust spatial working memory through homeostatic synaptic scaling in heterogeneous cortical networks.** *Neuron*, In press.
  34. Seung HS, Lee DD, Reis BY, Tank DW: **Stability of the memory of eye position in a recurrent network of conductance-based model neurons.** *Neuron* 2000, **26**:259-271.
  35. Koulakov AA, Raghavachari S, Kepecs A, Lisman JE: **Model for a robust neural integrator.** *Nat Neurosci* 2002, **5**:775-782.  
The first paper to point out, in the context of line attractors for neuroscience, that continuous variables could be closely and robustly approximated by a discrete representation. This study set out the basic ideas that were then also implemented, using alternative biological mechanisms, in [38] and [39].
  36. Lisman JE, Fellous JM, Wang XJ: **A role for NMDA-receptor channels in working memory.** *Nat Neurosci* 1998, **1**:273-275.
  37. Egorov AV, Hamam BN, Franssen E, Hasselmo ME, Alonso AA: **Graded persistent activity in entorhinal cortex neurons.** *Nature* 2002, **420**:173-178.  
A remarkable paper that describes the first *in vitro* model of graded persistent activity. With synaptic activity blocked and bath-applied carbachol (a cholinergic agonist), individual neurons of rat entorhinal cortex slices are capable of displaying graded persistent activity. Although the induction protocol used in this paper seems far from physiological, the mere fact that single *in vitro* neurons can display such activity opens important methodological doors for studying its mechanistic basis.
  38. Goldman MS, Levine JH, Major G, Aksay E, Tank DW, Seung HS: **Dendritic bistability increases the robustness of persistent neural activity in a model oculomotor neural integrator.** *Soc Neurosci Abstracts* 2002, **32**:266.14.
  39. Loewenstein Y, Sompolinsky H: **Computation by calcium dynamics in single neurons: a possible solution to the problem of neural integration.** *Soc Neurosci Abstracts* 2002, **32**:266.13.
  40. Lisman JE, Goldring MA: **Feasibility of long-term storage of graded information by the Ca<sup>2+</sup>/calmodulin-dependent protein-kinase molecules of the postsynaptic density.** *Proc Natl Acad Sci USA* 1988, **85**:5320-5324.
  41. Petersen CCH, Malenka RC, Nicoll RA, Hopfield JJ: **All-or-none potentiation at CA3-CA1 synapses.** *Proc Natl Acad Sci USA* 1998, **95**:4732-4737.
  42. Goldman MS, Maldonado P, Abbott LF: **Redundancy reduction and sustained firing with stochastic depressing synapses.** *J Neurosci* 2002, **22**:584-591.
  43. Alkon DL, Rasmussen H: **A spatial-temporal model of cell activation.** *Science* 1988, **239**:998-1005.
  44. Fuster JM: **Unit activity in prefrontal cortex during delayed-response performance: neuronal correlates of transient memory.** *J Neurophysiol* 1973, **36**:61-78.
  45. Kubota K, Niki H: **Prefrontal cortical unit activity and delayed alternation performance in monkeys.** *J Neurophysiol* 1971, **34**:337-347.
  46. Brody CD, Hernandez A, Zainos A, Lemus L, Romo R: **Timing and neural encoding of somatosensory parametric working memory in macaque prefrontal cortex.** *Cerebral cortex*, in press.  
The authors describe time-varying responses of delay-active prefrontal cortex neurons in detail, and show the 'time-stretching' of 'late' neuron responses when the delay period is lengthened.
  47. Rainer G, Miller EK: **Timecourse of object-related neural activity in the primate prefrontal cortex during a short-term memory task.** *Eur J Neurosci* 2002, **15**:1244-1254.  
This paper describes a strongly time-dependent delay period of persistent activity that occurred during a memory task that requires remembering one out of a discrete set of items. Interestingly, the temporal profile of the activity changes with experience. Like the current review, this paper argues that modelers ignore the very commonly observed time-varying nature of persistent activity at their peril.
  48. Leon MI, Shadlen MN: **Representation of elapsed time by neurons in area lip of the macaque.** *Soc Neurosci Abstracts* 2000, **30**:249.13.
  49. Leon MI, Shadlen MN: **Effect of expected reward magnitude on the response of neurons in the dorsolateral prefrontal cortex of the macaque.** *Neuron* 1999, **24**:415-425.
  50. Rainer G, Rao SC, Miller EK: **Prospective coding for objects in primate prefrontal cortex.** *J Neurosci* 1999, **19**:5493-5505.
  51. Fuster JM: **The prefrontal cortex — an update: time is of the essence.** *Neuron* 2001, **30**:319-333.  
A review that, like this one, argues that the role of time during delay periods is important.