

Touché: the feeling of choice

Peter E Latham & Peter Dayan

A new model by Machens *et al.* proposes a mechanism by which prefrontal cortex neurons can do two jobs that are normally thought to occur independently. In a stimulus comparison task, these model neurons both cast votes for a stimulus and make decisions.

“The people who cast the votes decide nothing. The people who count the votes decide everything.” —Stalin

To decide which flavor of ice cream to buy, you might taste first the mint chip and then the rocky road. To decide which painting to hang, you could try first the Picasso and then the Matisse. To decide which boy to date, you might kiss first Jack and then Sam. Decisions like these involving a comparison between memories and immediate sensations are a key part of our daily lives. Not surprisingly, a major goal of neuroscience is to understand the nature of decision-making and the workings of the neural structures underlying it. In a most interesting paper in a recent issue of *Science*, Machens and colleagues¹ have taken us one step closer to this goal. By proposing a circuit that implements both the memory and the comparison—one that both casts and counts votes—they have provided us with a parsimonious theory that ties up several experimental loose ends that have been vexing theoretical neuroscientists.

The paper investigates the mechanisms underlying a somatosensory version of a two-stimulus interval discrimination task, a basic workhorse of psychophysics and sensory physiology. In this task, subjects compare two successively presented stimuli, in this case two frequencies of vibration of a touch probe, and report the higher frequency. This task involves two primitive computations—memory for the first frequency (f_1) during the interstimulus interval, and comparison between the stored f_1 and the frequency of the second stimulus (f_2).

A seemingly obvious neural circuit for solving this problem (Fig. 1) has two components for the two computations: working memory for remembering the first stimulus, and a mutually inhibitory winner-take-all network for making the comparison and the eventual decision. It would seem reasonable to locate the working memory in prefrontal cortex (PFC), because PFC has long been recognized to be involved in working memory^{2,3} and is also known to receive information about vibration frequency (in this case f_1) from higher somatosensory cor-




Figure 1 Hypothetical circuit for comparing stimuli that occur at different times. **(a)** Loading phase. The first stimulus, f_1 , which is on in this period, is encoded in the firing rates of neurons in S2. This code is passed to PFC. **(b)** Maintenance phase. Persistent activity in PFC stores the value of f_1 during the delay period, which lasts for several seconds. During this phase, S2 is silent. **(c)** Comparison phase. Stimulus f_1 is sent to one population of neurons in the comparison circuit (which sits inside the rectangular box). Stimulus f_2 , which is active in this period, is sent to another population in the same circuit. These populations, which interact through mutual inhibition, form a winner-take-all network. The population with the higher frequency—and thus, in this model, the higher firing rate—ultimately silences the other population.

tex (S2)⁴, an area that encodes frequency during stimulus presentation. It is harder to pin down the comparison system, although based on data from the lateral intraparietal area (LIP) during decision-making, one might infer that structures associated with reporting could be involved. Unfortunately, this perfectly beautiful theory is ruined by inconvenient data. There are indeed anatomically relevant PFC neurons that code, via persistent activity, for f_1 during the interstimulus interval. However, these neurons come, rather unnecessarily, in two flavors: ‘plus’ neurons with monotonically increasing and decreasing tuning curves to f_1 and f_2 , respectively, and ‘minus’ neurons with the opposite tuning profile^{1,4}. More importantly, the activity of these same neurons seems to report which of the two stimuli the monkey chooses. Because these neurons report the monkey’s decision, it is plausible that they are actively involved in the decision-making circuit, thus casting doubt on the ‘obvious’ two-component model that postulates separate circuits for memory and comparison.

Machens *et al.*¹ report two insights that make limpid sense of these data. First, they note that the natural opposition of the plus and minus neurons is reminiscent of the kind of opposition that occurs in winner-take-all networks, exactly the structure assumed in the comparison network in Figure 1. Second, they observe that minor changes in the parameters

of winner-take-all networks can turn them into the sorts of memory devices (in this case a line attractor⁵) capable of storing quantities such as f_1 . Combining these two insights led them to a network that both supports memory and compares the two frequencies, thus making sense of the seemingly redundant plus and minus neurons, and also nicely matching various experimental data from PFC¹.

Figure 2 shows how the authors’ proposed mechanisms work. The basic network involves plus and minus groups of units, with mutual inhibition between the groups. By changing external control parameters, the network can be made to exist in three dynamical regimes. To switch the network among these regimes, Machens *et al.* use external input as a control parameter (an idea also widely used in Grossberg’s opponent dipole networks⁶); for didactic convenience in describing their mechanism, we will use self-excitation within each of the two groups as the control parameter.

For weak self-excitation, the network has a stable operating point that is determined by the input from S2. In this regime, the location of this stable point is a function of the first stimulus that is presented, f_1 . A high-frequency vibration (large f_1), for example, would cause the plus neurons to fire at high rate and the minus neurons at low rate (Fig. 2a); a low-frequency vibration would do the opposite.

Peter Latham and Peter Dayan are at the Gatsby Computational Neuroscience Unit, Alexandra House, 17 Queen Square, London WC1N 3AR, UK. e-mail: pel@gatsby.ucl.ac.uk

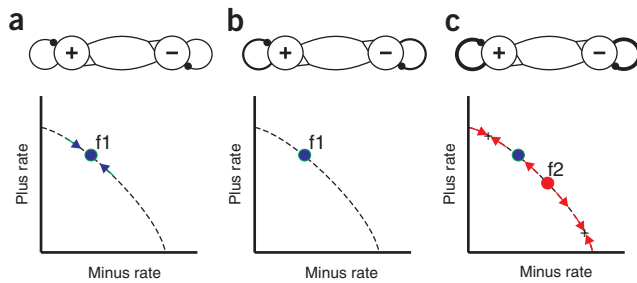


Figure 2 The Machens *et al.* mechanism. The top of each panel shows the mutual excitation circuit they use. Filled circles correspond to excitatory connections, with strength indicated by the thickness of the lines, and open triangles correspond to inhibitory connections. The bottom of each panel shows the firing rates of the plus (+) and minus (-) groups plotted against each other. Because of the mutual inhibition, there is a general trend for the rates to be anti-correlated; thus, the interesting dynamics happen along the dashed line. More specifically, independent of any external control parameter, the firing rates are pulled quickly onto the dashed line and after that evolve more slowly along the line. **(a)** Loading phase. Recurrent excitation is weak, so the firing rates of both the plus and minus neurons are set by f_1 . The blue arrows indicate flow along the line; all initial firing rates are attracted to the point indicated by the blue ball. **(b)** Maintenance phase. Recurrent excitation is slightly stronger, so there is no drift along the dashed line. In this regime, the firing rates of the two groups of neurons remain constant. **(c)** Comparison phase. Recurrent excitation is stronger still, and the circuit switches to winner-take-all. The threshold is set by the value of f_2 : firing rates above it are pushed toward the fixed point at the top of the dashed line (the plus between the red arrows); firing rates below it are pushed toward the fixed point at the bottom. In this example, f_1 is larger than f_2 , so the plus group ultimately fires at high rate and the minus group becomes relatively silent.

For somewhat stronger self-excitation, the network becomes neutrally stable, with activity tending to persist wherever it starts (**Fig. 2b**). In the hands of Machens *et al.*, this regime corresponds to the delay period, in which activity in PFC is approximately constant.

Finally, for yet stronger self-excitation, the network acts as a winner-take-all, with just one group dominating. In this regime, the stimulus input f_2 affects almost exclusively which state the network settles into. In effect, it sets a threshold relative to the current activity of the network, and thus biases which of the two populations will ultimately be active (**Fig. 2c**). For this mechanism to report the larger frequency correctly, f_2 must be coupled to the plus and minus groups in an opposite manner to f_1 , with higher f_2 boosting the minus group more than the plus group, and vice versa. Machens *et al.* avail themselves of some creative wiring to achieve this.

Machens *et al.* present a beautiful theory, but is it also ruined by inconvenient data? There are two interesting potential challenges to their theory. One comes from the temporal dynamics during the delay period. In the model, the firing rates of the PFC neurons are rock solid. Real prefrontal neurons, however, do not have such stable activity, but rather show a net change in firing rate (see Figures 1c,d and 4a of Machens *et al.*, which show delay data and delay model¹). Furthermore, we can expect from other studies that a significant fraction of the neurons in PFC (>90%) will show strong time dependence during the delay^{7,8}. Such changes are potentially informative of the expected time at which a

decision will be made (and this has been the target of some highly instructive analysis by the same group⁸). However, for the decision model, these modifications also make the point of comparison for the second stimulus potentially capricious. This effect can be ameliorated if both plus and minus neurons change firing rates in the same direction, indicating a drift perpendicular to the line attractor rather than parallel to it. Whether or not this is the case in real PFC neurons needs to be checked experimentally. Particularly problematic for the model would be an experimental finding of net drift parallel to the line attractor without bias in the monkeys' responses. This is because a key prediction is that position on the line attractor is the only source of memory.

The other potential challenge comes from a discrepancy between the temporal dynamics of the model and the real neural activity at the time of the second stimulus: the relatively simple dynamics of the model minus neurons at the time of the second stimulus are not matched by the data (again compare their Figures 1c,d and 4a). In the model, minus neurons show a monotonic decrease in firing rate when they report that the frequency of the second stimulus is higher than that of the first, and a non-monotonic convergence toward about 40 Hz in the opposite case. In the neural data, on the other hand, the minus neurons always show a transient increase in activity before a decision is made. How f_1 is preserved through this transient is not completely clear. Nor is it clear whether

this transient is a key element of the dynamics of PFC neurons that is beyond the scope of the model, or if it is an unimportant detail easily explained by minor modifications to it.

Independent of these challenges, Machens and colleagues have shown us how a single network can very naturally make a comparison between stimuli that occur at different times. The idea of having multiple different dynamics in the same underlying network is a familiar theme in discussions of the effects of neuromodulators on neural circuits⁹ and gated dipoles⁶, but is new in the current context. A network such as this is a critical element in building a theory of decision-making. It is furthermore attractively general purpose, as it will readily work with any one-dimensional input variable, not just simple stimuli like vibration frequency. In addition, given the ease with which line attractors can be made to perform integration over time⁵, the same network could handle situations in which evidence accumulates slowly, as it does whenever sensory input is noisy. The nature of the accumulated evidence is also flexible: it could be a function of the expected reward¹⁰ or of the log-likelihood ratio in favor of one or other alternatives (such as in motion discrimination tasks^{11,12}). In the latter case, the network could perform optimal statistical inference, an intriguing possibility in light of recent work on the relation between noisy integration and optimal decision-making¹³.

Decision-making lies toward the end stage of one of the most taxing problems organisms face: action selection in an uncertain world. It combines elements of optimal inference, game theory, memory, reward and punishment, and learning. Understanding the neural circuits that implement and integrate these elements is a key objective, toward which Machens *et al.* have taken us one step further.

1. Machens, C.K., Romo, R. & Brody, C.D. *Science* **307**, 1121–1124 (2005).
2. Fuster, J.M. & Alexander, G.E. *Science* **173**, 652–654 (1971).
3. Funahashi, S., Bruce, C.J. & Goldman-Rakic, P.S. *J. Neurophysiol.* **61**, 331–349 (1989).
4. Romo, R., Brody, C.D., Hernandez, A. & Lemus, L. *Nature* **399**, 470–473 (1999).
5. Seung, H.S. *Proc. Natl. Acad. Sci. USA* **93**, 13339–13344 (1996).
6. Grossberg S. *Biol. Psychiatry* **48**, 81–98 (2000).
7. Rainer, G. & Miller, E.K. *Eur. J. Neurosci.* **15**, 1244–1254 (2002).
8. Brody, C.D., Hernandez, A., Zainos, A. & Romo, R. *Cereb. Cortex* **13**, 1196–1207 (2003).
9. Marder, E. & Thirumalai, V. *Neural Netw.* **15**, 479–493 (2002).
10. Platt, M.L. & Glimcher, P.W. *Nature* **400**, 233–238 (1999).
11. Gold, J.I. & Shadlen, M.N. *J. Neurosci.* **23**, 632–651 (2003).
12. Mazurek, M.E., Roitman, J.D., Ditterich, J. & Shadlen, M.N. *Cereb. Cortex* **13**, 1257–1269 (2003).
13. Brown, E., Gao, J., Holmes, P., Bogacz, R., Gilzenrat, M. & Cohen, J.D. *Int. J. Bifurcation Chaos* (in the press).