



Flexible Control of Mutual Inhibition: A Neural Model of Two-Interval Discrimination

Christian K. Machens, *et al.*

Science **307**, 1121 (2005);

DOI: 10.1126/science.1104171

The following resources related to this article are available online at www.sciencemag.org (this information is current as of April 8, 2008):

Updated information and services, including high-resolution figures, can be found in the online version of this article at:

<http://www.sciencemag.org/cgi/content/full/307/5712/1121>

Supporting Online Material can be found at:

<http://www.sciencemag.org/cgi/content/full/307/5712/1121/DC1>

A list of selected additional articles on the Science Web sites **related to this article** can be found at:

<http://www.sciencemag.org/cgi/content/full/307/5712/1121#related-content>

This article **cites 29 articles**, 8 of which can be accessed for free:

<http://www.sciencemag.org/cgi/content/full/307/5712/1121#otherarticles>

This article has been **cited by** 39 article(s) on the ISI Web of Science.

This article has been **cited by** 19 articles hosted by HighWire Press; see:

<http://www.sciencemag.org/cgi/content/full/307/5712/1121#otherarticles>

This article appears in the following **subject collections**:

Neuroscience

<http://www.sciencemag.org/cgi/collection/neuroscience>

Information about obtaining **reprints** of this article or about obtaining **permission to reproduce this article** in whole or in part can be found at:

<http://www.sciencemag.org/about/permissions.dtl>

4. W. J. Gehring, M. G. H. Coles, D. E. Meyer, E. Donchin, *Psychophysiology* **27**, S34 (1990).
5. J. Hohnsbein, M. Falkenstein, J. Hoorman, *J. Psychophysiol.* **3**, 32 (1989).
6. P. S. Bernstein, M. K. Scheffers, M. G. H. Coles, *J. Exp. Psychol. Hum. Percept. Perform.* **21**, 1312 (1995).
7. S. Ito, V. Stuphorn, J. W. Brown, J. D. Schall, *Science* **302**, 120 (2003).
8. C. S. Carter *et al.*, *Science* **280**, 747 (1998).
9. M. M. Botvinick, L. Nystrom, K. Fissel, C. S. Carter, J. D. Cohen, *Nature* **402**, 179 (1999).
10. A. W. MacDonald, J. D. Cohen, V. A. Stenger, C. S. Carter, *Science* **288**, 1835 (2000).
11. C. B. Holroyd, M. G. Coles, *Psychol. Rev.* **109**, 679 (2002).
12. W. Schultz, P. Dayan, P. R. Montague, *Science* **275**, 1593 (1997).
13. J. O'Doherty, P. Dayan, K. Friston, H. Critchley, R. Dolan, *Neuron* **38**, 329 (2003).
14. J. W. Brown, D. Bullock, S. Grossberg, *Neural Networks* **17**, 471 (2004).
15. S. Bao, V. T. Chan, M. M. Merzenich, *Nature* **412**, 79 (2001).
16. Materials and methods are available as supporting material on Science Online.
17. G. D. Logan, W. B. Cowan, *Psychol. Rev.* **91**, 295 (1984).
18. K. Rubia *et al.*, *Neuroimage* **13**, 250 (2001).
19. V. Stuphorn, T. L. Taylor, J. D. Schall, *Nature* **408**, 857 (2000).
20. D. Boussaoud, S. P. Wise, *Exp. Brain Res.* **95**, 28 (1993).
21. M. M. Botvinick, T. S. Braver, D. M. Barch, C. S. Carter, J. C. Cohen, *Psychol. Rev.* **108**, 624 (2001).
22. A. D. Jones, R. Cho, L. E. Nystrom, J. D. Cohen, T. S. Braver, *Cogn. Affect. Behav. Neurosci.* **2**, 300 (2002).
23. K. R. Ridderinkhof, M. Ullsperger, E. A. Crone, S. Nieuwenhuis, *Science* **306**, 443 (2004).
24. In addition to the ACC effects, there was also a small but reliable difference in behavioral performance across the two go conditions, with reaction times slightly slower for high/go versus low/go [736 ms versus 710 ms, $F(1,15) = 5.80$, $P < 0.05$]. Because of this pattern, a potential alternative is that the reaction time differences are the causal mechanism underlying differences in ACC activity (if this ACC region was purely sensitive to response-related activity) rather than error-likelihood effects. To rule out this potential confound, ACC activity was reestimated after partially correlating out reaction time as a nuisance covariate. The effects of high/go and low/go remained significant in regions 1 [$t(14) = 2.67$, $P < 0.02$] and 2 [$t(14) = 2.29$, $P < 0.04$].
25. A whole-ACC ROI consisting of all Brodmann areas 24 and 32 voxels that showed significant change > go effects also showed effects of high/change > low/change ($P < 0.04$). The high/go > low/go effect did not reach significance in this ROI ($P = 0.21$), but high/go activity was numerically greater than low/go in 288 of 291 voxels (binomial test, $P < 0.00001$). High and low activities also apparently diverged with training as expected, but this did not reach significance.
26. R. Herrnstein, *J. Exp. Anal. Behav.* **4**, 267 (1961).
27. L. P. Sugrue, G. S. Corrado, W. T. Newsome, *Science* **304**, 1782 (2004).
28. C. B. Holroyd *et al.*, *Nat. Neurosci.* **7**, 497 (2004).
29. C. D. Fiorillo, P. N. Tobler, W. Schultz, *Science* **299**, 1898 (2003).
30. The authors thank C. Hoyer for help with data collection and C. Holroyd, J. Reynolds, A. Schaeffer, J. Schall, and V. Stuphorn for helpful comments. Supported by Office of Naval Research N00014-00-1-0715. Partial funding support was also received by the Director of Central Intelligence's Intelligence Technology Innovation Center. The authors declare that they have no competing financial interests.

Supporting Online Material
www.sciencemag.org/cgi/content/full/307/5712/1118/DC1

Materials and Methods
Fig. S1
Table S1

28 September 2004; accepted 27 December 2004
10.1126/science.1105783

Flexible Control of Mutual Inhibition: A Neural Model of Two-Interval Discrimination

Christian K. Machens,¹ Ranulfo Romo,² Carlos D. Brody^{1*}

Networks adapt to environmental demands by switching between distinct dynamical behaviors. The activity of frontal-lobe neurons during two-interval discrimination tasks is an example of these adaptable dynamics. Subjects first perceive a stimulus, then hold it in working memory, and finally make a decision by comparing it with a second stimulus. We present a simple mutual-inhibition network model that captures all three task phases within a single framework. The model integrates both working memory and decision making because its dynamical properties are easily controlled without changing its connectivity. Mutual inhibition between nonlinear units is a useful design motif for networks that must display multiple behaviors.

In our daily lives, our minds can flit from thought to thought with remarkable speed and flexibility (*1*). A simplified task that requires rapid shifts between different mental actions is known as two-interval discrimination (two stimuli separated by a time interval; Fig. 1A). Subjects must first perceive a brief stimulus, called *f1*, maintain it in working memory for several seconds, and then compare it with a brief second stimulus, called *f2*, to immediately decide which of the two stimuli was larger. The task requires both working memory and decision making, interfacing between the two in a rapid switch from one to the other.

The biophysical mechanisms underlying the performance of this task remain unknown.

Spiking neural-network models, built to serve as mechanistic accounts of cognitive neural activity, have focused so far on only single cognitive processes (*2–8*). Few models (*9, 10*), and no spiking network models, have addressed the question of how more than one computation and dynamic can be implemented in a single network. Yet cognitive acts typically require more than one type of computation. Many cognitive psychology models do integrate multiple processes, but do not address biophysical mechanisms (*11*). On the basis of recent neurophysiological data (Fig. 1) (*12–17*), we use a nonlinear dynamical systems approach (*18–21*) to design a simple and testable spiking-neuron model of two-interval discrimination. The model integrates three key processes into a single framework that proposes mechanistic links between the different processes, as well as between biophysical properties and neural and behavioral phenomena. These processes are fast initial loading of stimulus *f1* into working memory, slow

maintenance of working memory, and fast decision making.

Figure 1, C and D show the firing rates of two prefrontal cortical (PFC) neurons recorded from previously trained macaque monkeys while they performed a two-interval discrimination task in which *f1* and *f2* were the frequencies of mechanical vibrations applied to the tip of a finger (*12, 16, 22*). The dynamics of the activity of these neurons depends strongly on the phase of the task. During the loading of *f1* into working memory, there is a rapid flow to an *f1*-dependent firing rate. During the maintenance of *f1* in working memory, there is a long-lasting persistence of *f1*-dependent firing rates, despite the absence of the stimulus. During the comparison/decision phase, upon presentation of the second stimulus *f2*, the firing rates quickly segregate into one of two categories, depending on the monkey's subsequent choice of a "yes" or "no" push-button answer to the question, "Is *f1* greater than *f2*?" Responses similar to these PFC responses are also found in ventral (*17*) and medial (*14*) premotor cortices. For brevity, here we will refer collectively to these three areas as "frontal lobe areas." We highlight two aspects of the frontal lobe data. First, signals are encoded in complementary sets of roughly equal numbers of neurons (*12, 14, 17*). One set is composed of "plus" neurons, defined as neurons with a delay-period firing rate that is a monotonically increasing function of *f1* (Fig. 1C). Plus neurons typically fire the most for "yes" decisions after presentation of *f2*. The complementary set are "minus" neurons, defined as those which have delay period firing rates that are monotonically decreasing functions of *f1*, and fire the most for "no" decisions (Fig. 1D). Because higher *f2* stimuli are more likely to lead to "no" decisions; plus neurons are excited by high *f1* stimuli but inhibited by high *f2* stimuli. The converse occurs for minus neurons

¹Cold Spring Harbor Laboratory, 1 Bungtown Road, Cold Spring Harbor, NY 11724, USA. ²Instituto de Fisiología Celular, Universidad Nacional Autónoma de México, 04510 México, D.F., México.

*To whom correspondence should be addressed. E-mail: brody@csh.edu

(fig. S9) (23). Each of the two sets of neurons carries all the information necessary for the task; the existence of two apparently redundant sets is so far unexplained.

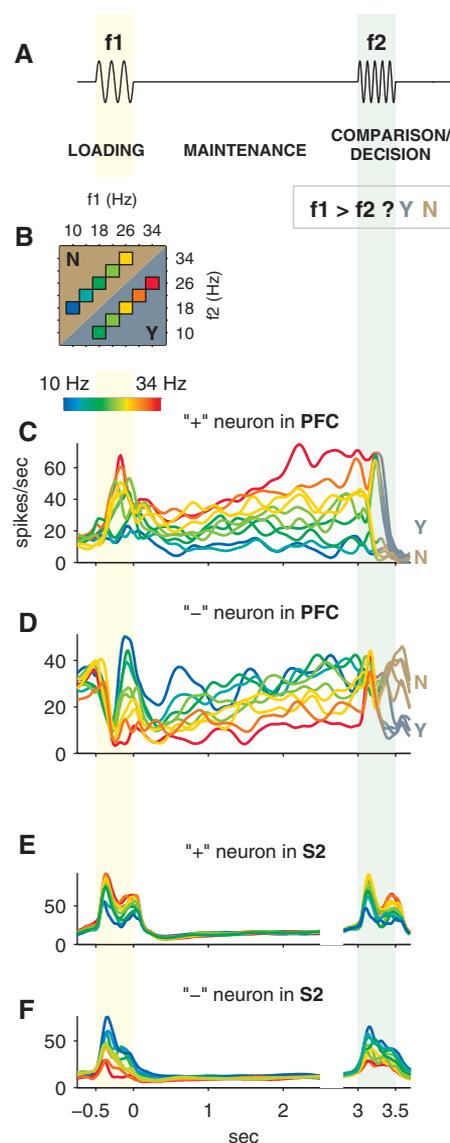


Fig. 1. The two-interval discrimination task and neuronal responses to it in the PFC and the S2 cortices. (A) Schematic diagram of each trial of the task. (B) Typical stimulus set used in the neurophysiological studies we focus on here (14, 17, 44, 45). Stimuli are frequencies of mechanical vibrations applied to the tip of a monkey's finger, and each colored box indicates a (f1, f2) stimulus pair. For each pair, monkeys made the correct response more than 91% of the time (16, 22). (C to F) Neuronal responses. The rainbow color code at the upper left indicates the f1 value applied during each type of trial. Y/N color code indicates the push button pressed by the monkey at the end of each trial. (C) and (D) show smoothed firing rates of two different PFC neurons recorded over many trials. (C) shows a plus neuron and (D) shows a minus neuron. (E) and (F) show neuronal responses in area S2 (13, 15). For (E) and (F) only, rainbow colors for $t > 2.8$ s indicate value of f2.

Second, the same neurons that show graded delay-period activity, which represents working memory of f1, also show categorical activity after stimulus f2, which represents the monkey's decision (14, 17, 24). This finding contrasts with current mathematical psychology models of two-interval discrimination, which have implicitly assumed that working memory and decision making are processes represented by separate variables (25–28). Instead, we propose an algorithm (Fig. 2) in which both memory and decision outcome are represented by the value of a single state variable (horizontal axis). The dynamical modes of the system are described by a hypothetical energy function L (vertical axis), the shape of which does not depend on the value of the state variable. The state variable always evolves so as to reduce L . During the loading phase, the external stimulus creates a single minimum in L (a single stable point) at a position determined by the value of f1. This forces the state into an f1-dependent position. During the memory maintenance phase, there is no longer an external stimulus that determines the shape of L . The memory of f1 is represented by the state's position, and for this to remain steady, the L function must be approximately flat [a line attractor configuration (2, 29, 30)]. During the comparison/decision phase, we map stimulus f2 onto the same horizontal axis as f1. All state positions to the left of f2 now represent memories of f1 values that are less than f2. States to the right of f2 represent memories of f1 that are greater than f2. If a peak in L (an unstable point) is created at the position given by f2, the state will evolve in one of two opposite directions, depending on the yes or no answer to $f1 > f2$? If the horizontal axis stands for a firing rate that grows from left to right, the plots mimic the activity of plus neurons through all phases of the task. If the firing rate grows from right to left, the plots mimic minus neurons.

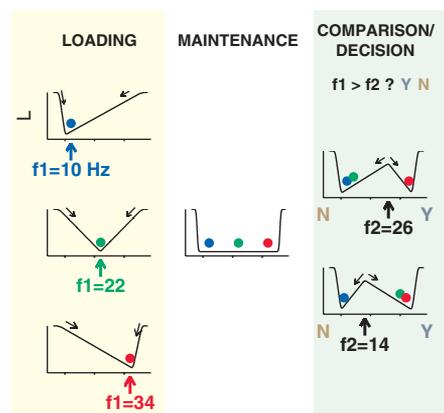


Fig. 2. One-dimensional dynamical algorithm for two-stimulus-interval discrimination. The value of the state variable (horizontal axis) is used to represent both memory of f1 (maintenance) and decision-making outcome (comparison/decision).

We propose that the frontal-lobe areas instantiate this dynamical algorithm. The task-relevant sensory inputs to these frontal areas arise from the secondary somatosensory cortex (S2) (Fig. 1, E and F) (31). During the first stimulus, responses in the S2 are similar to those in the PFC: there are both plus (Fig. 1E) and minus (Fig. 1F) neurons (13, 15). Unlike the PFC, however, during most of the delay period, S2 neuron firing rates are low and not f1-dependent. Also in contrast with the PFC, immediately after presentation of f2, neurons in the S2 respond to f2 with the same plus- or minus-type firing-rate dependence with which they responded to f1 (32).

Clues about the underlying frontal-lobe neural architecture come from the analysis of firing rate covariations between pairs of PFC neurons. These tend to be positive if both neurons are plus, or if both are minus, but negative when one is plus and one is minus (33), leading us to consider the architecture sketched in Fig. 3A. In a simplified version of this circuit (Fig. 3B), each node represents a population of neurons, and each node's output variable is the average activity of the population. Figure 3C shows a node's i/o function, which is defined as its output, expressed as a resulting postsynaptic conductance, as a function of excitatory input (34). The overall shape obtained, with a threshold below which output is negligible and with saturation at high outputs, is characteristic of many neuron models.

To study graphically the dynamics of the circuit in Fig. 3B, we show in black in Fig. 3D the output of the plus node as a function of the inhibitory input from the minus node. An additional excitatory input, E , is held constant here. The minus node's i/o function can be plotted by exchanging the horizontal and vertical axes to form the brown axes and curve in Fig. 3D. This phase-plane plot now describes the complete dynamics of the system, because we can follow the input to output activity of each node as it reverberates around the circuit loop. Points where the two i/o curves intersect are known as fixed points, or steady states. These may be stable (like the minimum in L during loading as seen in Fig. 2) or unstable (like the maximum in L during comparison/decision).

During the first stimulus [loading phase (Fig. 3E)], inputs from the S2 area are active. Together with the external input E , they can shift the i/o functions along their input axes. If both nodes receive the same S2 input, then by symmetry, the crossing point of the two i/o curves must lie along the 45° diagonal of the phase-plane plots. However, we propose that S2 plus neurons project to frontal plus neurons, and S2 minus neurons project to frontal minus neurons. As a result, the position of the single stable point is determined by the value of f1, thus instantiating the loading mode of Fig. 2. Horizontal position in Fig. 2 corre-

sponds here to the angle of θ (defined in top phase plane, Fig. 3E).

During the delay period [maintenance phase (Fig. 3F)], the inputs from S2 are silent. The value of E and the neuronal-model parameters can be chosen such that the two i/o curves largely overlap. This creates a quasi-continuous line of stable points (a line attractor) and thus instantiates the maintenance mode of Fig. 2 (35).

During the comparison/decision phase (Fig. 3G), we propose that an external control signal, which indicates that the current stimulus should be treated as f2, not f1, reduces the otherwise constant excitatory input E . For low enough E , the system has two stable fixed points on either side of a single unstable fixed point, as required for this phase in Fig. 2. In the data, frontal neurons switch the sign of their stimulus dependence between f1 and f2 (23), but S2 sensory neurons do not. The net functional connection between S2 and frontal neurons must therefore change sign. To match this in the model, we used the circuit shown in Fig. 3H to switch the net S2-to-frontal plus neuron connections between loading (Fig. 3E, top) and comparison/decision (Fig. 3G, top). A similar circuit was used for frontal minus neuron connections (36). As a result of this input sign switch, increasing f2 moves the unstable fixed point from lower to higher θ (Fig. 3G), which matches the f1-dependence of the stable fixed point during loading (Fig. 3E). This completes instantiation of the comparison/decision mode of Fig. 2.

The continuous-variable nodes of the model of Fig. 3 were each replaced by 250 noisy, leaky, integrate-and-fire neurons in order to produce a spiking neuron model with almost identical behavior (fig. S4). Figure 4A shows firing rates for one spiking neuron from the plus node and one neuron from the minus node, qualitatively capturing key aspects of the data. The sign of correlations between pairs of neurons in the model also matches the pattern found in the experimental data (Fig. 4B).

The one-dimensional algorithm of Fig. 2 produces testable predictions separate from the neural instantiation proposed here. If the decision is reported when the state reaches one of the two final stable points in Fig. 2, and if the distance to the final stable point is a dominant factor in the time required to reach it, then it follows that on trials where f1 equals f2 (50% “yes” responses, 50% “no”) and f1 and f2 are near the high end of the range of stimuli, “yes” decisions will be reached faster than “no” decisions. The converse is true at the low end of the range (Fig. 4D). At the neural instantiation level, both the fine-tuned (Fig. 3) and robust (34, 35) instantiations rest on populations of mutually inhibitory neurons. If the strength of inhibitory connections were increased [e.g., by use of benzodiazepines (37)],

then the memory maintenance mode would be perturbed in the direction of the comparison/decision mode, with a single unstable point. Because of the circuit’s symmetry, the effect would be similar to an f2 stimulus in the middle of the range (≈ 22 Hz), even though no

stimulus would be present (Fig. 4C). The psychophysical correlate would be a tendency to categorize f1 into high or low rather than remembering it accurately.

Unlike previous line attractor models (3, 5, 6), which relied mostly on mutual exci-

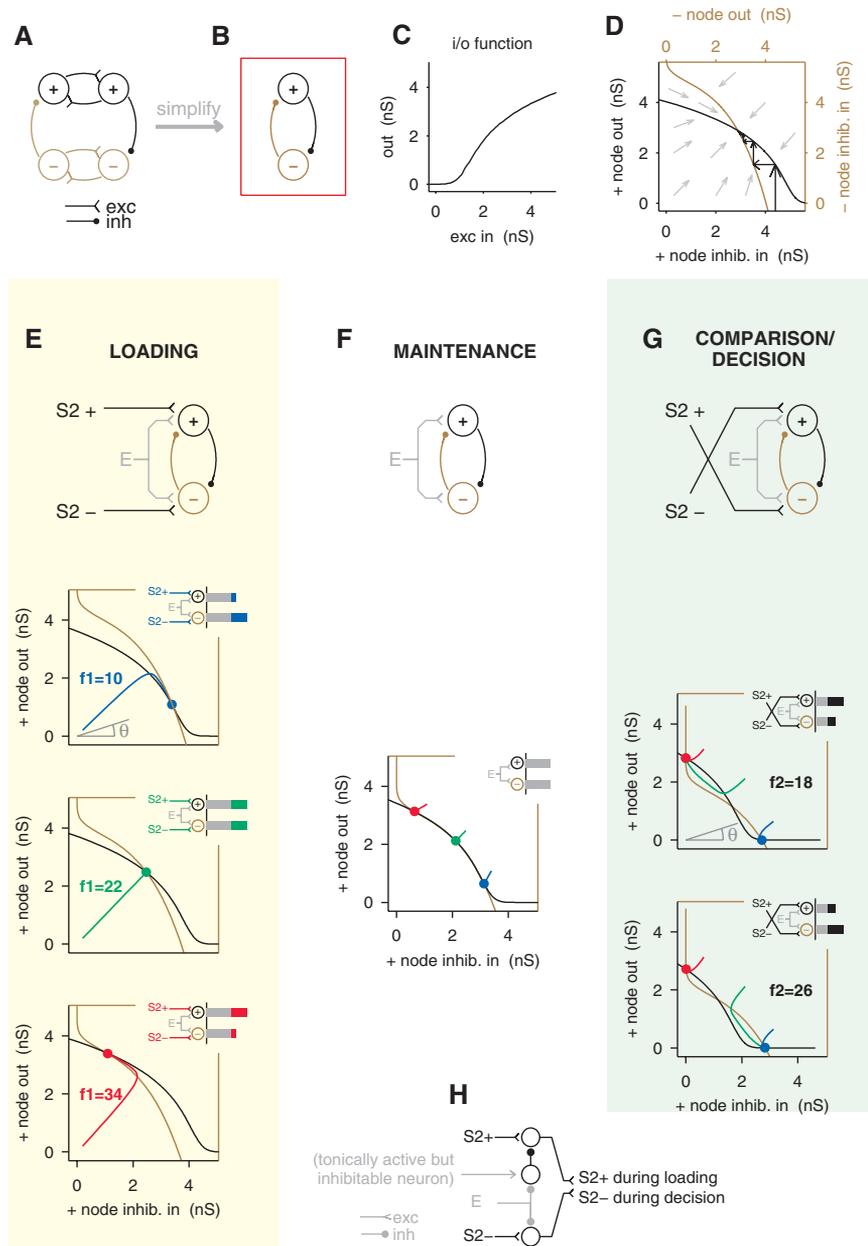


Fig. 3. Mutual inhibition model implements the dynamical modes of Fig. 2. (A and B) Circuit diagrams. Each node represents a population of neurons. (C) Input/output of a node as a function of excitatory input (nS, nanoSiemens). (D) Phase-plane plot of both i/o functions allows tracing out the dynamics of the mutual inhibition circuit. (Black arrows indicate stepping discretely between nodes, for illustration, and gray arrows indicate continuous time dynamics, as used in actual simulations.) (E) Loading. The top panel shows inputs to the circuit. The lower three panels show the effect of three different values (10, 22, and 34 Hz) of the first stimulus (f1). Colored lines show dynamical trajectories, starting from (0,0) and ending at filled circles. (F) Maintenance. Removing inputs from S2 results in a quasi-continuum of stable points. Dynamical trajectories start from correspondingly colored endpoints of trajectories in (E). (G) Comparison/decision. The lower panels show i/o functions and dynamics for f2 = 18 (top) and f2 = 26 (bottom). The final state depends on the answer to f1 > f2? Trajectories start from the endpoints of (F). For all plots in (E), (F), and (G), angle θ corresponds to horizontal position in Fig. 2. (H) Input switching circuitry. At high E , S2 plus signals pass through, but at low E , S2 minus signals pass through.

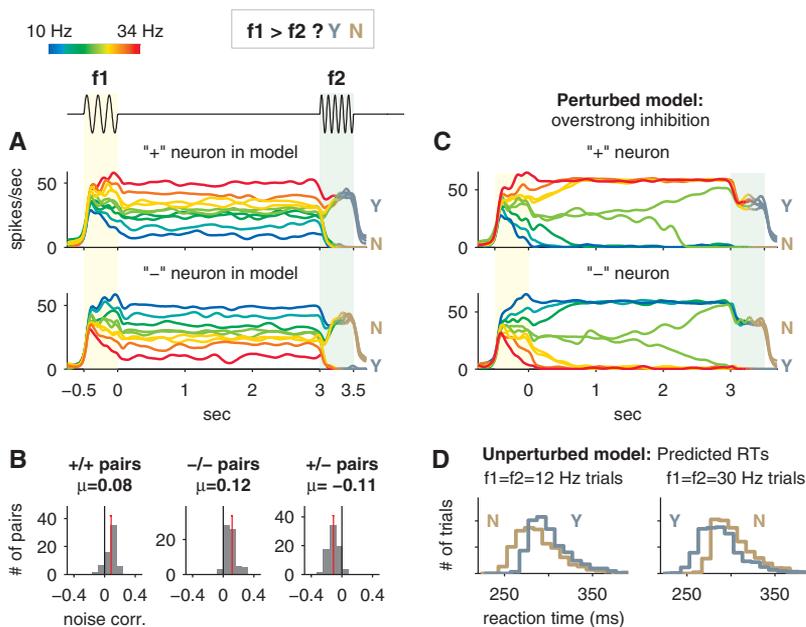


Fig. 4. Model responses and predictions. **(A)** Same format and stimulus values as Fig. 1, C and D. Simulated activity for one neuron from the plus node population (top) and one neuron from the minus node population (bottom), averaged over trials. **(B)** Histograms of noise correlations (41) during the maintenance phase, compiled over pairs of neurons in the model. **(C)** Predicted firing rates in model where inhibition is overstrong. **(D)** One-dimensional algorithm (Fig. 2) predicts asymmetry in “yes” and “no” reaction times (RTs). See text for sufficient conditions under which this prediction follows. Particular RT values shown here were obtained by using the model of Fig. 3 with small amounts of Gaussian noise added at each time step.

tation, we implemented both working-memory and decision-making dynamics through mutual inhibition. This facilitated the use of a single simple circuit for both. The use of neuron models with i/o functions that are nonlinear above threshold, instead of linear or threshold-linear models (38–40), further allowed straightforward control of the circuit’s dynamical mode through an external excitation signal.

Our daily mental lives have an enormous variety of highly flexible dynamics. What is the neural basis of this flexibility? Do frontal lobes contain many separate modules of neurons, each capable of a particular type of computation and its attendant dynamics? Or, as the data of Fig. 1, C and D have inspired us to propose here, can single modules of frontal lobe neurons rapidly reconfigure their dynamical properties, switching between different behaviors as the cognitive flow requires? If single modules are indeed flexible, what is the range of dynamics and computations that they can display? We have only begun to address these questions here. But we believe they are fundamental, and lie at the heart of the nature of the neural architecture underlying cognition.

References and Notes

1. E. K. Miller, J. D. Cohen, *Annu. Rev. Neurosci.* **24**, 167 (2001).
2. H. S. Seung, *Proc. Natl. Acad. Sci. U.S.A.* **93**, 13339 (1996).
3. H. S. Seung, D. D. Lee, B. Y. Reis, D. W. Tank, *Neuron* **26**, 259 (2000).
4. X. J. Wang, *Neuron* **36**, 955 (2002).

5. P. Miller, C. D. Brody, R. Romo, X. J. Wang, *Cereb. Cortex* **13**, 1208 (2003).
6. A. A. Koulakov, S. Raghavachari, A. Kepecs, J. E. Lisman, *Nature Neurosci.* **5**, 775 (2002).
7. Y. Loewenstein, H. Sompolinsky, *Nature Neurosci.* **6**, 961 (2003).
8. E. Brown et al., *Int. J. Bifurcat. Chaos*, in press.
9. S. L. Moody, S. P. Wise, G. di Pellegrino, D. Zipser, *J. Neurosci.* **18**, 399 (1998).
10. E. Salinas, *Neural Comput.* **15**, 1439 (2003).
11. A. Miyake, P. Shah, Eds., *Models of Working Memory: Mechanisms of Active Maintenance and Executive Control* (Cambridge Univ. Press, Cambridge, ed. 1, 1999).
12. R. Romo, C. D. Brody, A. Hernandez, L. Lemus, *Nature* **399**, 470 (1999).
13. E. Salinas, A. Hernandez, A. Zainos, R. Romo, *J. Neurosci.* **20**, 5503 (2000).
14. A. Hernandez, A. Zainos, R. Romo, *Neuron* **33**, 959 (2002).
15. R. Romo, A. Hernandez, A. Zainos, L. Lemus, C. D. Brody, *Nature Neurosci.* **5**, 1217 (2002).
16. C. D. Brody, A. Hernandez, A. Zainos, R. Romo, *Cereb. Cortex* **13**, 1196 (2003).
17. R. Romo, A. Hernandez, A. Zainos, *Neuron* **41**, 165 (2004).
18. H. R. Wilson, *Spikes, Decisions, Actions: Dynamic Foundations of Neuroscience* (Oxford Univ. Press, Oxford, ed. 1, 1999).
19. H. D. Abarbanel, M. I. Rabinovich, *Curr. Opin. Neurobiol.* **11**, 423 (2001).
20. P. Dayan, L. F. Abbott, *Theoretical Neuroscience* (MIT Press, Cambridge, MA, ed. 1, 2001).
21. F. C. Hoppensteadt, *An Introduction to the Mathematics of Neurons* (Cambridge Univ. Press, Cambridge, ed. 2, 1997), pp. 155–158.
22. A. Hernandez, E. Salinas, R. Garcia, R. Romo, *J. Neurosci.* **17**, 6391 (1997).
23. PFC neurons change the net sign of their stimulus dependence between stimuli f1 and f2. That is, the higher the stimulus value applied during the loading phase, the higher a PFC plus neuron’s firing rate and the lower the minus neuron’s firing rate. Yet the higher the stimulus value applied during comparison/decision, the lower the PFC plus neuron’s firing rate and the higher the PFC minus neuron’s firing rate

because higher f2s are more likely to lead to “no” decisions (fig. S9).

24. C. Brody, et al., in preparation.
25. D. Green, J. Swets, *Signal Detection Theory and Psychophysics* (J. Wiley, New York, 1966).
26. R. Ratcliff, *Psychol. Rev.* **88**, 552 (1981).
27. S. Link, *The Wave Theory of Difference and Similarity* (L. Erlbaum, Hillsdale, NJ, 1992).
28. R. Luce, *Response Times: Their Role in Inferring Elementary Mental Organization* (Oxford Univ. Press, Oxford, 1991).
29. X. J. Wang, *Trends Neurosci.* **24**, 455 (2001).
30. C. D. Brody, R. Romo, A. Kepecs, *Curr. Opin. Neurobiol.* **13**, 204 (2003).
31. S. T. Carmichael, J. L. Price, *J. Comp. Neurol.* **363**, 642 (1995).
32. During the late part of the second stimulus, the firing of many S2 neurons becomes correlated with the monkey’s decision (15), but this occurs after decision-correlated activity arises in the PFC, and thus cannot be causal to it. For simplicity, in the model we treat only the initial creation of decision-correlated activity in the PFC, before it has propagated back to area S2.
33. During the delay period and across trials using identical f1 stimuli, we calculated the noise correlation r between pairs of neurons (41). We found covariation between pairs of plus neurons ($r = 0.13 \pm 0.04$ (SE), $n = 32$ pairs), and pairs of minus neurons ($r = 0.08 \pm 0.03$, $n = 17$). But pairs where one neuron is plus and the other is minus were anticorrelated ($r = -0.08 \pm 0.02$, $n = 58$) (24).
34. Materials, methods, and documented computer code are available as supporting material on Science Online.
35. Good overlap in Fig. 3F requires high precision in the design of the i/o function shapes. This is a generic difficulty that all line attractors face (30, 42). We address this issue in the supporting online text by adapting the solution proposed by Koulakov et al. (6) to create a more robust model. Despite its greatly increased dimensionality, the more robust model can still be approximately described in terms of two variables (fig. S6), similar to the diagrams of Fig. 3.
36. The circuit of Fig. 3H is only one possible way to switch input signs. See (43) for another possibility.
37. U. Rudolph, H. Mohler, *Annu. Rev. Pharmacol. Toxicol.* **44**, 475 (2004).
38. A. Morishita, A. Yajima, *Kybernetik* **11**, 154 (1972).
39. H. Seung, in *The Handbook of Brain Theory and Neural Networks*, M. Arbib, Ed. (MIT Press, Cambridge, MA, ed. 2, 2003), pp. 94–97.
40. M. Usher, J. L. McClelland, *Psychol. Rev.* **108**, 550 (2001).
41. E. Zohary, M. N. Shadlen, W. T. Newsome, *Nature* **370**, 140 (1994).
42. H. S. Seung, D. D. Lee, B. Y. Reis, D. W. Tank, *J. Comput. Neurosci.* **9**, 171 (2000).
43. E. Salinas, *J. Neurosci.* **24**, 1113 (2004).
44. R. Romo, E. Salinas, *Annu. Rev. Neurosci.* **24**, 107 (2001).
45. R. Romo, E. Salinas, *Nature Rev. Neurosci.* **4**, 203 (2003).
46. C.K.M. and C.D.B. jointly developed and implemented the model and wrote the paper. R.R. conceived and carried out the experiments. This work was supported in part by a Swartz Foundation Fellowship to C.K.M.; an International Research Scholars Award from the Howard Hughes Medical Institute and awards from Dirección General de Asuntos del Personal Académico–Universidad Autónoma Nacional de México and the Millennium Science Initiative–Consejo Nacional de Ciencia y Tecnología to R.R.; and by a Sloan Foundation Research Fellowship, an award from the Redwood Neuroscience Institute, and NIH grant 1R01MH067991-01 to C.D.B. We thank J. D. Cohen, Z. F. Mainen, S. S. -H. Wang, and A. M. Zador for comments on the manuscript, and J. J. Hopfield for discussion.

Supporting Online Material

www.sciencemag.org/cgi/content/full/307/5712/1121/DC1
 Materials and Methods
 Figs. S1 to S9
 References
 Model S1

17 August 2004; accepted 6 January 2005
 10.1126/science.1104171