

Modeling Working Memory and Decision Making Using Generic Neural Microcircuits^{*}

Prashant Joshi

Institute for Theoretical Computer Science
Technische Universität Graz
A-8010 Graz, Austria
joshi@igi.tugraz.at
<http://www.igi.tugraz.at/joshi>

Abstract. Classical behavioral experiments to study working memory typically involve three phases. First the subject receives a stimulus, then holds it in the working memory, and finally makes a decision by comparing it with another stimulus. A neurocomputational model using generic neural microcircuits with feedback is presented here that integrates the three computational stages into a single unified framework. The architecture is tested using the two-interval discrimination and delayed-match-to-sample experimental paradigms as benchmarks.

1 Introduction

Classical experiments in behavioral and experimental neuroscience that are employed to study the working memory typically involve three phases. First the subject receives a stimulus, then this stimulus is held in working memory, and finally a decision is made by comparing it with another incoming stimulus. Two classical experimental paradigms to test this are the two-interval discrimination [1] and the delayed-match-to-sample [2] tasks.

A recent study proposed controlled mutual inhibition of neurons in pre-frontal cortex (PFC) as the neural algorithm behind the working memory and decision making process [1]. Although the model was successful in integrating the loading, maintenance, and decision making phases, obtaining such precise tuning of mutual inhibition in cortical circuits in PFC seems biologically implausible. Also, despite existing evidence that shows synaptic learning as a responsible mechanism for working memory related tasks [2], the model used a static (no learning involved) neural circuit.

This article demonstrates that simple linear readouts from generic neural microcircuit models that send their output as a feedback signal to the circuit, can be used to model decision making process that involves the use of working memory. The neurocomputational architecture described here integrates the three crucial stages described above into a single unified framework. These sequential

^{*} The work was partially supported by the Austrian Science Fund FWF, project #P17229 – N04 and project #FP6 – 015879 (FACETS) of the European Union.

stages are the initial loading (L) of stimulus into working memory, maintenance (M) of working memory, and decision (D) making (following the notation used in [1]). Essentially this model presents a unified computational framework for working memory and decision making carried out by the neurons in PFC. For comparison, the unified framework is used to build a spiking neural network model of two-interval discrimination. Additionally to show that this computational paradigm is task-independent, we employ the same paradigm to model the delayed-match-to-sample task.

The core principles behind the working of this model make the assumption that the cortex can be thought of as an ultra-high dimensional dynamical system, where the afferent inputs arriving from thalamus and the recurrent cortical feedbacks are churned in a non-linear way to obtain a high-dimensional projection of the low-dimensional input space. Preceding work has demonstrated that such high dimensional transient dynamics provides the neural circuit with analog fading memory¹ that provides the circuit enough computational power for performing open-loop sensory processing tasks [3, 4].

Analog fading memory by itself is not powerful enough to render the circuits the power to hold information in working memory. The obvious reason being that analog fading memory by itself has an upper limit on the order of tens of msec, depending on the time constants of synapses and neurons in the neural circuit [4], whereas typically the working memory holds information on the order of seconds. Recent results show that feedback from trained readout neurons that are part of generic neural circuit can induce multiple co-existing “partial attractors” in the circuit dynamics [5, 6]. This result is further extended here to demonstrate that even in the presence of feedback noise, such “partial attractor” states can be held by generic neural circuits on the time-scales of several seconds, that is obviously a requirement for tasks involving working memory.

The neural microcircuit model considered in this article is described in section 2, and the results for the two-interval discrimination and delayed-match-to-sample tasks are presented in sections 3 and 4 respectively. Finally section 5 presents a discussion of the results.

2 Generic Neural Microcircuit Models

For the experiments described in this article, generic cortical microcircuit models consisting of integrate-and-fire neurons were used, with a high level of noise that reflects experimental data. Biologically realistic models of dynamic synapses were used whose individual mixture of pair-pulsed depression and facilitation

¹ A map (or filter) F from input- to output streams is defined to have *fading memory* if its current output at time t depends (up to some precision ε) only on values of the input \mathbf{u} during some finite time interval $[t - T, t]$. Formally, F has fading memory if there exists for every $\varepsilon > 0$ some $\delta > 0$ and $T > 0$ so that $|(F\mathbf{u})(t) - (F\tilde{\mathbf{u}})(t)| < \varepsilon$ for any $t \in \mathbb{R}$ and any input functions $\mathbf{u}, \tilde{\mathbf{u}}$ with $\|\mathbf{u}(\tau) - \tilde{\mathbf{u}}(\tau)\| < \delta$ for all $\tau \in [t - T, t]$. This is a characteristic property of all filters that can be approximated by an integral over the input stream \mathbf{u} , or more generally by Volterra- or Wiener series.

(depending on the type of pre- and postsynaptic neuron) was based on experimental data [7, 8]. These circuits were not created for any specific computational task. Sparse synaptic connectivity between neurons was generated (with a biologically realistic bias towards short-range connections) by a probabilistic rule², and synaptic parameters were chosen randomly from distributions that depended on the type of pre- and postsynaptic neurons (in accordance with empirical data from [7, 8])³. The neurons in the generic cortical microcircuit models were placed on the integer-points of a 3-D grid⁴, and 20% of these neurons were randomly chosen to be inhibitory.

Each readout neuron was trained by linear regression to output at any time t , a particular target value $f(t)$. Linear regression was applied to a set of data points of the form $\langle y(t), f(t) \rangle$, for many time points t , where $y(t)$ is the output of low-pass filters applied to the spike-trains of pre-synaptic neurons, and $f(t)$ is the target output. During training, the feedback from readouts performing diverse computational tasks was replaced by a *noisy* version of their target output (“teacher forcing”) ⁵. Note that teacher forcing with noisy versions of target feedback values trains these readouts to correct errors resulting from imprecision in their preceding feedback (rather than amplifying errors)[9].

The generic neural microcircuit model received analog input streams from 4 sources for the experiment modeling two-interval discrimination (from 5 sources in the experiment modeling delayed-match-to-sample). The outcomes of the experiments discussed in this article were all negative if these analog input streams were directly fed into the circuit (as input current for selected neurons in the circuit). Apparently the variance of the resulting spike trains were too large to

² The probability of a synaptic connection from neuron a to neuron b (as well as that of a synaptic connection from neuron b to neuron a) was defined as $C \cdot \exp(-D^2(a, b)/\lambda^2)$, where $D(a, b)$ is the Euclidean distance between neurons a and b , and λ is a parameter which controls both the average number of connections and the average distance between neurons that are synaptically connected. Depending on whether the pre- or postsynaptic neuron were excitatory (E) or inhibitory (I), the value of C was set according to [8] to 0.3 (EE), 0.2 (EI), 0.4 (IE), 0.1 (II). For the experiment modeling a) the two-interval discrimination task, $C = 1$, $\lambda = 1.5$; b) the delayed-match-to-sample task, $C = 1$, $\lambda = 1.2$.

³ *Neuron Parameters*: membrane time constant 30 ms, absolute refractory period 3 ms (excitatory neurons), 2 ms (inhibitory neurons), threshold 15 mV (for a resting membrane potential assumed to be 0), reset voltage drawn uniformly from the interval [13.8, 14.5] mV, constant non-specific background current I_b uniformly drawn from the interval [13.5, 14.5] nA for each neuron, noise at each time-step I_{noise} drawn from a gaussian distribution with mean 0 and SD chosen for each neuron randomly from a Gaussian distribution over the interval [4.0, 5.0] nA, input resistance $1 M\Omega$, the initial condition for each neuron, i.e. its membrane potential at time $t = 0$, was drawn randomly (uniform distribution) from the interval [13.5, 14.9] mV.

⁴ 400(500) neurons, arranged on a $20 \times 5 \times 4$ ($20 \times 5 \times 5$) grid for the circuit modeling the two-interval discrimination (delayed-match-to-sample) task.

⁵ At each time-step t , a different noise value of $0.0001 \times \rho \times f(t)$ was added, where ρ is a random number drawn from a gaussian distribution with mean 0 and SD 1, and $f(t)$ is the current value of the input signal (signal-dependent noise).

make the information about the slowly varying values of these input streams readily accessible to the circuit. Therefore the input streams were instead fed into the circuit with a simple form of spatial coding, where the location of neurons that were activated encoded the current values of the input variables. More precisely, each input variable is first scaled into the range $[0, 1]$. This range is linearly mapped onto an array of 50 symbolic input neurons. At each time step, one of these 50 neurons, whose number $n(t) \in \{1, \dots, 50\}$ reflects the current value $i_n(t) \in [0, 1]$ which is the normalized value of input variable $i(t)$ (e.g. $n(t) = 1$ if $i_n(t) = 0$, $n(t) = 50$ if $i_n(t) = 1$). The neuron $n(t)$ then outputs at time step t , the value of $i(t)$. In addition the 3 closest neighbors on both sides of neuron $n(t)$ in this linear array get activated at time t by a scaled down amount according to a gaussian function (the neuron number n outputs at time step t the value $i(t) \frac{1}{\sigma\sqrt{2\pi}} \exp \frac{-(n-n(t))^2}{2\sigma^2}$, where $\sigma = 0.8$). Thus the value of each of the $4(5)^6$ input variables is encoded at any time by the output values of the associated 50 symbolic input neurons (of which at least 43 neurons output at any time the value 0). The neuron in each of these $4(5)$ linear arrays are connected with one of the $4(5)$ layers consisting of 100 neurons in the previously described circuit of $((20 \times 5) \times 4)$ ($((20 \times 5) \times 5)$) integrate-and-fire neurons.

The coding scheme used here is similar to the population coding [10], i.e., the neurons in the circuit fire maximally for their own preferred stimuli, and different neurons have different preferred stimuli. This coding scheme is suitable in context of this article as previous studies demonstrate that neurons in PFC process information via a labelled-line code [11, 12].

3 Two-Interval Discrimination

In the actual experimental protocol for two-interval discrimination task, a pre-trained subject is presented with two frequencies f_1 and f_2 , separated by a certain delay interval. Initially the frequency f_1 is loaded into the working memory, and its value is maintained during the delay phase, and on presenting the f_2 frequency, the subject is required to decide whether " $f_1 > f_2$?". Two kind of neurons have been observed in PFC which show opposite activities in response to the above question [1]. The first type of neurons (called "+" neurons from now), show an increase in their activity during the D phase, when the answer to the above question is "yes", and the second type of neurons (called "-" neurons from now), show an increase in their neural activity when the answer to the above question is "no". The information required to carry out the task is present in the firing activity of "+" and "-" neurons independently, and the reason for the simultaneous presence of both these sets is till now unknown.

In this setup⁷, the "+" and "-" neurons have been modeled as simple linear readouts that send feedback of their activity to the neural circuit (see Figure 1).

⁶ 4 input variables for the two-interval discrimination task, and 5 input variables for the delayed-match-to-sample task.

⁷ Total simulation time for one trial 3.5 s, simulation time-step 10 ms, f_1 and f_2 are presented for 0.5 s each, during the L and D phases respectively.

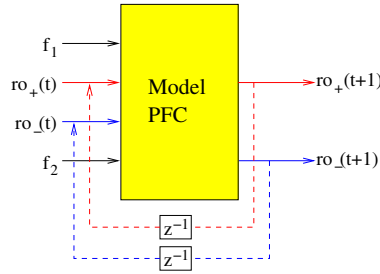


Fig. 1. Closed-loop setup for the two-interval discrimination task. The model PFC circuit is made of 400 integrate-and-fire neurons arranged on the integer points of a $20 \times 5 \times 4$ cube. The circuit receives 2 frequencies (f_1 , f_2) as external inputs, and two feedback signals ($ro_+(t)$, $ro_-(t)$) from the “+” and “-” readouts. The task is to answer the question if “ $f_1 > f_2$?” The “+” (“-”) neurons show an increase in their activity when the answer to the above question is “yes” (“no”). The notation z^{-1} denotes a unit time-step delay in feedback signal.

In addition to the feedback, the circuit receives 2 external inputs (f_1 and f_2). The input signals f_1 and f_2 are presented during the L and D phases respectively (see Figure 2 A). The training data consisted of 2 noisy versions of each of the 10 input pairs (f_1 , f_2) (see Figure 2 B). The target functions of “+” and “-” readouts are as described above. Figure 2 E shows a 200 ms blowup of the circuit response of 100 randomly chosen neurons (activity of excitatory neurons shown in black) during one of the closed-loop validation runs ($f_1 = 18$ Hz, $f_2 = 26$ Hz). The panels C and D of figure 2 show the desired (black) and observed values of the “+” and “-” readouts during this run. Panel F and G show the response of the “+” and “-” readouts for the 10 pairs of input frequencies, (f_1 , f_2) (note the similarity to Figure 1, panels C and D in [1], which show the actual data from “+” and “-” neurons in PFC during the two-interval discrimination task).

To test the robustness of the neural model, experiments were done where after the readouts have been trained, a subset κ_n (κ_n progressively increased from 0.5% to 5% in 10 trials such that $\kappa_n \subset \kappa_{n+1}$) of synapses converging onto the “+” readouts were randomly chosen and pruned (synaptic weight set to 0). The resulting setup was tested with the 10 frequency pairs (f_1 , f_2). Panels H and I of figure 2 show the result of these robustness tests. Panel H shows the mean and standard error of correlation values for progressively higher levels of pruned synapses. Panel I shows the resulting matrix of correlation values, where each square shows the correlation value for a particular validation run, for a particular pruning percentage. The control correlation values⁸ (no pruning) are shown in the row on the top. Results indicate that the model shows graceful degradation in presence of suboptimal inputs. This is quite interesting, as traditional models of attractor based neural computation fail to demonstrate robustness [13].

⁸ Over 10 validation runs, mean = 0.9556, SD = 0.0206.

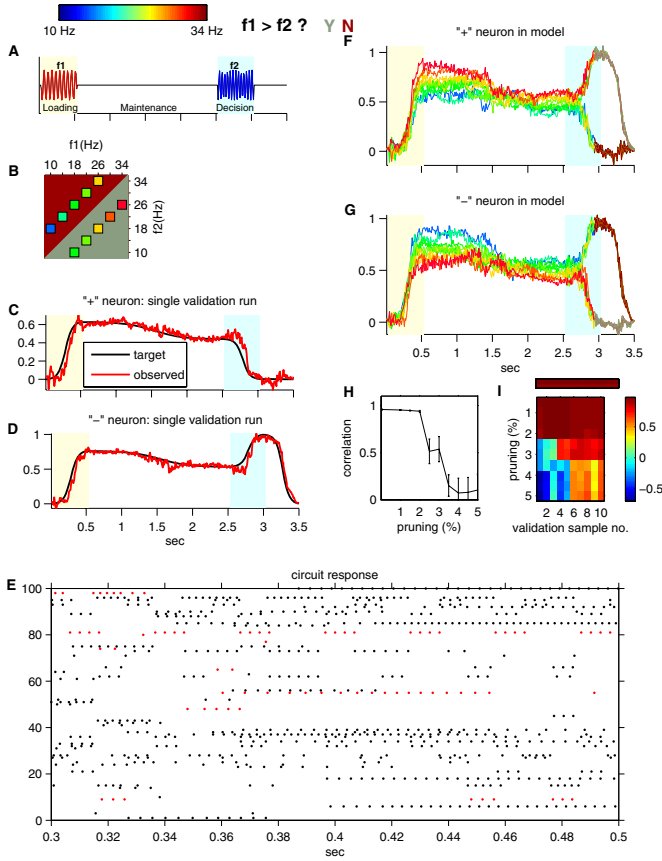


Fig. 2. (A) The external frequencies f_1 (18 Hz), and f_2 (26 Hz) presented during the L and D phase for one of the close-loop validation trials. (B) Stimulus frequencies used in this study. The target (black) and observed values for the (C) “+” readout, and (D) the “-” readout. (E) A blowup of 200 ms of resulting firing activity of 100 randomly chosen neurons in the circuit. Excitatory neurons are shown in black. Responses of the (F) “+” and (G) “-” readouts for each of the frequency pairs. The colorbar at upper left indicates the value of f_1 used in each trial. (H) Mean and standard error of correlation values for trials with progressively higher pruning percentage. (I) The resulting matrix of correlation values where each square shows the correlation value for a particular validation run, for a particular pruning percentage. The control correlation values (no pruning) are shown in the row on the top.

4 Delayed-Match-to-Sample

In this task, a pre-trained subject is presented with a cue visual stimulus during the L phase (for example a small colored square on a predetermined region of the screen), which is followed by a delay period (M phase), and subsequently in the D phase two simultaneous probe stimuli are presented at two different

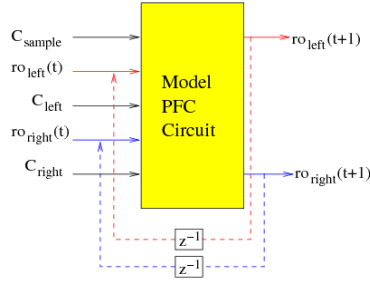


Fig. 3. Closed-loop setup for the delayed-match-to-sample task. The model PFC circuit is made of 500 integrate-and-fire neurons arranged on the integer points of a $20 \times 5 \times 5$ cube. The circuit received 3 color stimulus (C_{cue} , C_{left} , C_{right}) as external inputs, and 2 feedback signals ($ro_{left}(t)$, $ro_{right}(t)$) from the “left” and “right” readouts. The “left” (“right”) neurons show an increase in their activity during D phase when $C_{cue} = C_{left}$ (C_{right}).

places on the screen. The task is to decide which of the probe stimuli has the same color as the sample stimulus[2].

Figure 3 shows the setup⁹ used in this experiment. The model PFC circuit received 3 external inputs (C_{cue} , the cue color; C_{left} , the left probe color; C_{right} , the right probe color), and 2 feedback signals (from the “left” and “right” readouts). The cue stimulus was presented during the L phase and the probe stimuli were presented during the D phase. The neurons in the model PFC circuit made projections to two linear readouts (called “left” and “right” readouts from now on) which had similar functionality as the “+” and “-” readouts in the two-interval discrimination task. The “left” readout showed high activity if the answer to the question “ $C_{cue} = C_{left}$?” was “yes”. The “right” readout behaved exactly opposite to this and showed high amount of activity if the probe stimulus shown on the right matched the sample stimulus.

For this experiment the training data consisted of 2 noisy versions of each of the 5 input color triplets (C_{cue} , C_{left} , C_{right}) (see Figure 4 B). Figure 4 shows the result of a closed loop validation run. Panel A shows one such external triplet. The target (black) and observed response of the “left” and “right” readout are shown in panel C and D respectively. The panels E and F of figure 4 show the response of the “left” and “right” readouts for the 5 input stimuli (each line drawn in the color of corresponding cue stimulus).

To test the performance of the setup, we tested the setup in 100 different validation runs. Panels G and H of figure 4 present the histogram of correlation values for the “left”¹⁰ and “right”¹¹ readouts over the set of these 100 validation runs.

⁹ Total simulation time for one trial 1.95 s, simulation time-step 10 ms, C_{cue} is presented for 0.32 s during the L phase, and C_{left} and C_{right} are presented simultaneously for 0.6 s, during the D phase.

¹⁰ mean = 0.9962, SD = 0.0018.

¹¹ mean = 0.9942, SD = 0.0018.

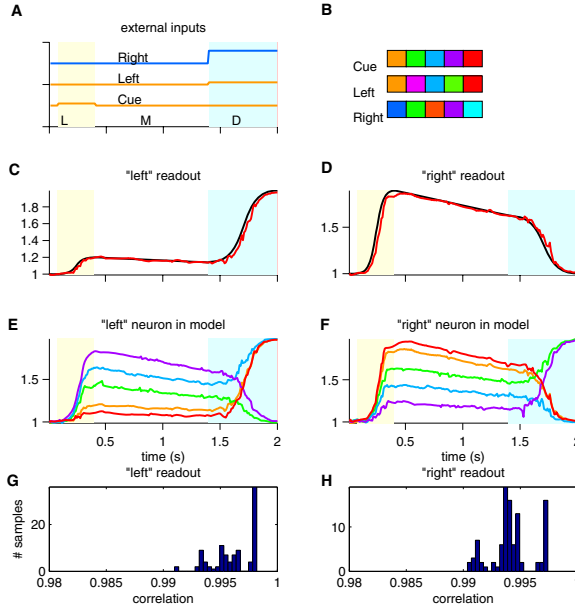


Fig. 4. (A) The external stimuli to the model PFC circuit during one of the closed loop validation runs. The cue stimulus is presented during the L phase and the left-and-right probe stimuli are presented during the D phase. (B) The colors of the cue, left, and right stimuli used in trials. The target (black) and observed values for the (C) “left” readout, and (D) the “right” readout. Responses of the (E) “left” and (F) “right” readouts for each of the color triplets. The performance of the setup was tested using 100 different closed loop validation runs. Histogram of correlation values for the (G) “left” and the (H) “right” readout.

5 Discussion

A new neurocomputational paradigm is described that uses synaptic learning mechanisms to present a unified model for working memory and decision making using biologically realistic neural microcircuit models composed of spiking neurons and dynamic synapses. Additionally results for the two-interval-discrimination task show that the neural algorithm is task independent. It is to be noted however that although spiking neural circuits were used to model the interval-discrimination tasks for added biological realism, it is not a requirement, and any recurrent circuit would give similar results, as long as it has the kernel property.

Readouts make a binary decision by reaching one of the states corresponding to the decision made by them. The actual point of time when the readout makes a decision can be thought of as a threshold crossing event, i.e. the first time when the readout crosses a threshold after the presentation of the probe stimulus in the D phase.

It was found that using population coding to encode the inputs projecting on to model PFC circuits was essential to obtain the demonstrated results. Using a population of neurons to encode each analog input stream is not unrealistic, as sufficient evidence exists in current literature of its existence. It is however not claimed here that the precise mechanism of population coding used in this article is the one used in cortical circuitry of PFC.

The feedback from the trained readouts played an apparently important role in the neural model described above. Although the generic neural microcircuits used to simulate the model PFC circuit are endowed with fading memory due to the kernel property of the circuits, this is not sufficient for holding information in working memory for longer timespans ranging in the order of seconds. Apparently the feedback from trained readouts provides the circuit with additional needed information that falls outside the window of fading memory, hence enhancing the information present in the circuit dynamics.

Obviously closed-loop applications of generic neural microcircuit models like the ones discussed in this article present a harder computational challenge than open-loop sensory processing tasks, since small imprecisions in their output are likely to be amplified by the plant (e.g. the arm model) to yield even larger deviations in the feedback, which is likely to further enhance the imprecision of subsequent outputs. This problem can be solved by teaching the readouts from the neural microcircuit during training to ignore smaller recent deviations reported by feedback, thereby making the target trajectory of output torques an attractor in the resulting closed-loop dynamical system.

This study also demonstrates the ability of generic neural microcircuit models to hold “partial attractor” states in their circuit dynamics for significantly longer and biologically relevant time-scales ranging in the order of a couple of seconds, in presence of noise. Also a point of interest is the robustness of this neurocomputational model to factors such as synaptic pruning, and feedback noise.

According to the “internal model” hypothesis [14, 15], there exists an internal model for each of the tasks that we have learned throughout our lives. One outcome of such a hypothesis would be that a given neuron may participate with a different synaptic weight in a number of neural assemblies, each supporting a different internal model. Interestingly, this is reflected in our setup too, as neurons in the generic neural circuit make synaptic projections to the set of readouts with different synaptic weights assigned for each task.

The results presented in this article demonstrate the role of feedback in enhancing the inherent fading memory of a neural circuit. Further it also shows the ability of generic neural circuits to model working memory and decision making, which happens at significantly longer time-scales. Further work is needed to explore if this neurocomputational architecture is extensible across diverse cognitive modalities, e.g. decision making, and motor control.

Acknowledgements

The author is thankful to Wolfgang Maass for critical comments on the draft version of the manuscript, and anonymous reviewers for helpful suggestions.

References

1. C. K. Machens, R. Romo, and C. D. Brody. Flexible control of mutual inhibition: A neural model of two-interval discrimination. *Science*, 307:1121–1124, 2005.
2. G. Rainer, H. Lee, and N. K. Logothetis. The effect of learning on the function of monkey extrastriate visual cortex. *PLoS Biology*, 2(2):275–284, 2004.
3. D. V. Buonomano and M. M. Merzenich. Temporal information transformed into a spatial code by a neural network with realistic properties. *Science*, 267:1028–1030, Feb. 1995.
4. W. Maass, T. Natschläger, and H. Markram. Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural Computation*, 14(11):2531–2560, 2002.
5. W. Maass, P. Joshi, and E. D. Sontag. Principles of real-time computing with feedback applied to cortical microcircuit models. *NIPS*, 2005.
6. W. Maass, P. Joshi, and E. D. Sontag. Computational aspects of feedback in neural circuits. *submitted for publication*, 2005.
7. H. Markram, Y. Wang, and M. Tsodyks. Differential signaling via the same axon of neocortical pyramidal neurons. *PNAS*, 95:5323–5328, 1998.
8. A. Gupta, Y. Wang, and H. Markram. Organizing principles for a diversity of GABAergic interneurons and synapses in the neocortex. *Science*, 287:273–278, 2000.
9. P. Joshi and W. Maass. Movement generation with circuits of spiking neurons. *Neural Computation*, 17(8):1715–1738, 2005.
10. A. Pouget and P. E. Latham. Population codes. In M. A. Arbib, editor, *The handbook of brain theory and neural networks*, pages 893–897. MIT Press, 2003.
11. J. M. Fuster. Unit activity in prefrontal cortex during delayed-response performance: neuronal correlates of transient memory. *J. Neurophysiol.*, 36:61–78, 1973.
12. J. M. Fuster, M. Bodner, and J. K. Kroger. Cross-modal and cross-temporal association in neurons of frontal cortex. *Nature*, 405:347–351, 2000.
13. H. S. Seung, D. D. Lee, B. Y. Reis, and D. W. Tank. Stability of the memory of eye position in a recurrent network of conductance-based model neurons. *Neuron*, 26(1):259–271, 2000.
14. Kenneth J. W. Craik. *The Nature of Explanation*. Cambridge University Press, 1943.
15. E. Bizzi and F. A. Mussa-Ivaldi. Neural basis of motor control and its cognitive implications. *Trends in Cognitive Sciences*, 2(3):97–102, 1998.