

Gain Mechanisms for Contextually Guided Visuomotor Transformations

Marina Brozović,^{1*} Alexander Gail,^{1,2*} and Richard A. Andersen¹

¹Division of Biology, California Institute of Technology, Pasadena, California 91125, and ²Bernstein Center for Computational Neuroscience, German Primate Center, 37077 Göttingen, Germany

A prevailing question in sensorimotor research is the integration of sensory signals with abstract behavioral rules (contexts) and how this results in decisions about motor actions. We used neural network models to study how context-specific visuomotor remapping may depend on the functional connectivity among multiple layers. Networks were trained to perform different rotational visuomotor associations, depending on the stimulus color (a nonspatial context signal). In network I, the context signal was propagated forward through the network (bottom-up), whereas in network II, it was propagated backwards (top-down). During the presentation of the visual cue stimulus, both networks integrate the context with the sensory information via a mechanism similar to the classic gain field. The recurrence in the networks hidden layers allowed a simulation of the multimodal integration over time. Network I learned to perform the proper visuomotor transformations based on a context-modulated memory of the visual cue in its hidden layer activity. In network II, a brief visual response, which was driven by the sensory input, is quickly replaced by a context-modulated motor-goal representation in the hidden layer. This happens because of a dominant feedback signal from the output layer that first conveys context information, and then, after the disappearance of the visual cue, conveys motor goal information. We also show that the origin of the context information is not necessarily closely tied to the top-down feedback. However, we suggest that the predominance of motor-goal representations found in the parietal cortex during context-specific movement planning might be the consequence of strong top-down feedback originating from within the parietal lobe or from the frontal lobe.

Key words: context; rule-guided behavior; sensorimotor transformations; recurrent networks; parietal cortex; feedback

Introduction

Depending on task rules, or more generally on the behavioral context, we can perform various motor actions with an object. In its simplest form, the task may correspond to a visuomotor remapping. Here, based on the position of a visual object and the behavioral rule, we can reach or turn our gaze toward or away from this object. In its more complex form, the task may involve throwing a ball, catching a ball, spinning it or rolling it, all depending on the rules of a game. Sensorimotor processing streams (Kalaska, 1996; Wise et al., 1997; Battaglia-Mayer et al., 2000; Buneo and Andersen, 2006) in the cerebral cortex must be flexible enough to produce radically different motor plans based on contextual information despite identical sensory conditions (Petrides, 1982; Platt and Glimcher, 1999; Wallis and Miller, 2003; Pasupathy and Miller, 2005).

Performing context-specific sensorimotor transformations means that nonspatial context information dynamically guides

the spatial remapping of a topographic sensory input (e.g., representing visual space) onto a topographic output (e.g., motor goal in extrinsic coordinates). In this study, we investigated how this task can be achieved via a gain-modulation process in simple, recurrent neural networks. Two recent theoretical studies addressed contextual visuomotor remapping. Salinas (Salinas, 2004) proposed that there exists a basis layer of stimuli-specific neurons that are gain modulated by the context information. Here, motor neurons read out the activation of the basis layer by direct linear addition. As opposed to our network, the basis network did not simulate sensory-context integration as a function of time. Furthermore, our network architecture allowed us to study how the feedback from the motor neurons affects the neural responses in the hidden layer. Another recent model addressed the related problem of selecting among multiple potential movement targets based on contextual information (Cisek, 2006). Among other differences to our present study, this previous model had the connections between the layers hardwired, and its complex architecture did not allow easy reconfiguration of the connection weights to study the effects of feedback within the network.

The contextual gain modulation in our case occurs as a direct result of the network learning process. Our models are a logical extension of the previous theoretical (Zipser and Andersen, 1988; Pouget and Sejnowski, 1995; Salinas and Abbott, 1996; Salinas and Thier, 2000) and neurophysiological (Brotchie et al., 1995;

Received Nov. 28, 2006; revised Aug. 4, 2007; accepted Aug. 14, 2007.

This work was supported by the Swartz Fellowship, Federal Ministry for Education and Research (Bundesministerium für Bildung und Forschung, Germany) Grant 01GQ0433, and the National Institutes of Health. We thank T. Yao and V. Shcherbatyuk for administrative and technical support.

*M.B. and A.G. contributed equally to this work.

Correspondence should be addressed to Dr. Marina Brozović, Mail Code 216-76, Division of Biology, California Institute of Technology, Pasadena, CA 91125. E-mail: brozovic@vis.caltech.edu.

DOI:10.1523/JNEUROSCI.2685-07.2007

Copyright © 2007 Society for Neuroscience 0270-6474/07/2710588-09\$15.00/0

Shenoy et al., 1999; Andersen and Buneo, 2003) studies of multimodal integrations. The proposed networks are different in that they investigate feedback and recurrent architectures, and they simulate the integration of qualitatively different information, namely sensory information with an abstract transformation rule. The first model assumed that the context information is processed via the same pathway as the sensory input. The network architecture is purely feed-forward with recurrent connections only within the hidden layer. The second model assumed that the context influences sensorimotor units indirectly via feedback connections originating from motor-goal units. In addition to investigating how different origins of the context signal shape the network behavior, we argue for the significance of the feedback signal based on anatomical-physiological grounds.

Materials and Methods

Visuomotor task

The visuomotor task is to perform a spatial remapping (clockwise rotation) of the visual cue position onto a motor goal position, depending on the color of the central fixation point (see Fig. 1A). This task may be implemented either in a reach or a saccade experiment (similar to the study by Takeda and Funahashi, 2002). The remapping rule defined in this study is viewed as a simplified version of more general contextual remapping, and thus the words “context” and “remapping rule” are used interchangeably throughout the article.

The timeline of the task is divided into three time periods with different durations. In the first period (and lasting throughout the trial), the central fixation point is highlighted in one of four colors conveying the mapping rule (blue, 0; green, 45; red, 90; yellow, 180° clockwise rotation). In the second period, the position of the visual cue is briefly flashed on a peripheral location in the visual field. The third period is the “memory” period, in which the subject prepares to execute a motor action such as a reach or saccade toward the remapped position of the visual cue. Note that in the two extreme mapping conditions (0 and 180°), this task corresponds to pro- and anti-saccades/reaches.

Network design

Proposed network models are of the three-layer recurrent network design with input, hidden, and output layers similar to the study by Xing and Andersen (2000b). The models have a one-dimensional (1D) space representation for the sensory input and the motor output, mimicking cue/motor-goal directional tuning as typically found in neurons of sensorimotor areas during center-out tasks.

There are eight Gaussian units encoding the position of the visual input and eight Gaussian units encoding the position of the motor goal in the output. The input tuning curves $R_i(\phi)$ have their centers ϕ_i^0 uniformly spaced from -180 to $+180^\circ$ in 45° intervals with σ also equal to 45° .

$$R_i = \exp\left(\frac{\cos(\varphi - \varphi_i^0) - 1}{\sigma^2}\right) \quad (1)$$

Note that these units cover the 1D circle in a periodic manner. The second input carries the information about the remapping rule (context), encoded by a single unit $C(\omega)$. Four discrete activation values (0.25, 0.5, 0.75, and 1) correspond to four desired rotation angles ω (90, 0, 180, and 45°). Alternatively, instead of a single $C(\omega)$ unit, we could use four Gaussian units, with partially overlapping peaks at 90, 0, 180, and 45° . This way, the task rule would be encoded as a “spatial” parameter (rotational angle) similar to the visual and motor information. The pooled activity of these contextual units would again act as a “gain” signal in each hidden unit (Zipsler and Andersen, 1988; Pouget and Sejnowski, 1995; Salinas and Abbott, 1996; Salinas, 2004). Hence, there is effectively no difference between the two implementations.

The desired activation of the output unit T_k is defined as follows:

$$T_k = \exp\left(\frac{\cos(\varphi - \varphi_k^0 - \omega) - 1}{\sigma^2}\right). \quad (2)$$

The goal of the network training was to learn the $R_i \rightarrow T_k$ mapping according to the context C . There were two layers of nonlinear transformations to achieve this. The network models discussed here had 40 hidden units, although we trained the networks with 30, 50, and 60 hidden units with qualitatively same results.

Network I

The idea of the network I architecture is for a direct spatial and context cue integration in a sensorimotor stage (hidden layer). Both types of input, the position and the context (e.g., color) cue, are originally visual sensory stimuli that in combination could be mapped onto corresponding motor goals. As argued in the Discussion, the context input to the sensorimotor stage could, in principle, be coming directly from sensory areas or be mediated via higher cognitive areas. Essential for network I is the parallel feedforward nature of the two inputs. Consequently, the first network design (see Fig. 1B) has a predominantly feedforward architecture, where the hidden layer receives inputs directly from R_i and C . The output layer units O_k receive inputs only from the hidden layer units H_j . The responses of the hidden units also depend on the previous activity of the layer via recurrent connections as follows:

$$H_j(t) = f\left(\sum_i w_{ji}^R R_i(t) + w_j^C C(t) + \sum_m w_{jm}^H H_m(t-1)\right) \quad (3)$$

$$O_k(t) = f\left(\sum_j w_{kj}^O H_j(t)\right).$$

Here, w^R denotes weights between the sensory input units R_i and hidden units H_j , w^{H} are the recurrent weights between current $H_j(t)$ and previous $H_m(t-1)$ activations of hidden units, and w^C are the weights connecting the single “rule” unit C with each of the units in the hidden layer. Finally, there is a set of weights w^O connecting hidden and output units. The function f denotes a transfer function used to obtain the activation in the hidden and output units. We used the sigmoid transfer function

$$f(\text{net}) = \frac{1}{1 + \exp(-\text{net})},$$

where net refers to the sum of the weighted input.

The learning rule for network I

We used the backpropagation through time algorithm (BPTT) to train the network (Werbos, 1990; Jaeger, 2002). The network weights were updated sequentially after each trial, which consisted of a single visual target and rule presentation. The training set contained 120 pairs of a position cue and a remapping rule. They were randomly selected from the total set of 72×4 possible input combinations (72 positions between -180 and $+180$ in increments of 5° ; four remapping rules). The progress of the network training was tested with another 120 randomly selected pairs. The network trained until the sum of squares error between the desired output T_k and the network output O_k dropped below 0.01. This implied that the fitted amplitudes and σ for each of the eight output Gaussians were within 10% of their ideal values. To achieve this level of performance required $\sim 70,000$ iterations.

The network was trained according to the timeline of events as described above and shown at the bottom of Figure 1B. In the first two time steps, the output layer contains only the information about the remapping rule C , the position of the visual cue V is not known yet. In the third time step, the hidden layer receives inputs from both V and C , as well as the “memory input” about its previous activity r_H [referring to all $H_j(t-1)$]. From this time step on, the output layer is required to produce the value of the motor goal m (i.e., the remapped visual cue position). In the following time intervals (four through eight), the memory about V is contained exclusively in the recurrent activity of the hidden layer units. The weights were initialized to some small random (positive and negative) values, which are updated according to the customized BPTT rules:

$$\begin{aligned}
\delta_k^{out}(t) &= (T_k(t) - O_k(t)) f'(O_k(t)) \\
\delta_j^{hh}(t_{max}) &= f'(H_j(t_{max})) \sum_k w_{jk}^O \delta_k^{out}(t_{max}) \\
\delta_j^{hh}(t-1) &= \left(\sum_k w_{jk}^O \delta_k^{out}(t-1) + \sum_m w_{jm}^{hh} \delta_m^{hh}(t) \right) f'(H_j(t-1)) \\
\Delta w_{jk}^{out} &= \sum_{t=t_{max}}^1 \delta_k^{out}(t) H_j(t), w_{jk}^{out} = w_{jk}^{out} + \eta \Delta w_{jk}^{out} \\
\Delta w_{jm}^{hh} &= \sum_{t=t_{max}}^2 \delta_j^{hh}(t) H_m(t-1), w_{jm}^{hh} = w_{jm}^{hh} + \eta \Delta w_{jm}^{hh} \\
\Delta w_{ij}^R &= \delta_j^{hh}(t=3) R_i(t=3), w_{ij}^R = w_{ij}^R + \eta \Delta w_{ij}^R \\
\Delta w_j^C &= \sum_{t=t_{max}}^1 \delta_j^{hh}(t) C(t), w_j^C = w_j^C + \eta \Delta w_j^C. \quad (4)
\end{aligned}$$

Here, the learning rate η is set to 0.01 and f' refers to:

$$\frac{df(\text{net})}{d\text{net}}.$$

The weight updates are calculated by first running the network activity forward in time ($t = 1 \dots 8$) and by keeping the key values such as $\delta_k^{out}(t)$, $H_j(t)$, and $O_k(t)$. This is followed by sequentially calculating weight changes in all time intervals, starting from the last one ($t_{max} = 8$) and moving to the beginning.

Network II

The idea of the second network is that the mapping rule information is qualitatively different from the purely sensory position information. It is only available to the sensorimotor areas after the color cue information has been evaluated by higher cognitive areas [e.g., in prefrontal cortex (PFC)], and a corresponding motor goal has been defined in the motor stages. Consequently, the hidden layer receives a direct input only from the Gaussian units R_i , which encode the position of the visual cue. The information about the C enters directly into the activation of the output units O_k via the weights w_k^C . As a result, the hidden units receive information about the remapping rule only indirectly, via the feedback weights w_{kl}^{FB} that reciprocally connect the output with the hidden layer. The hidden layer also receives its recurrent activity via w_{jm}^{hh} . In summary, the network is described by the following two equations:

$$\begin{aligned}
H_j(t) &= f \left(\sum_i w_{ji}^R R_i(t) + \sum_m w_{jm}^H H_m(t-1) + \sum_i w_{ji}^{FB} O_i(t-1) \right) \\
O_k(t) &= f \left(\sum_j w_{kj}^O H_j(t) + w_k^C C \right). \quad (5)
\end{aligned}$$

The learning rule for network II

The timeline of the network training is outlined at the bottom of Figure 1C, and it is identical to Network I. In the first time step, there is no input to the hidden layer units. This results in a uniform layer activation, because the transfer function is sigmoid:

$$f(0) = \frac{1}{1 + \exp(-0)} = 0.5.$$

Alternatively, it can be initialized to a small random number. The output units, however, receive a direct input from the rule unit C , as well as the weighted (w^O) input from the uniformly activated hidden layer. At this point, the output layer encodes only C . At $t = 2$, the rule information reaches the hidden layer not in its original form as it enters the output layer, but through the weighted (w^{FB}) activation of the output units from $t = 1$. The network II output layer still represents only rule C . At $t = 3$, the visual information V enters directly into the hidden layer and there it is combined with the indirect rule information as well as the recurrent activity of the hidden layer from the previous time step. The output units

continue to receive direct input from the rule unit C and the input from the hidden layer throughout the trial. Starting from $t \geq 3$, the output units of network II encode strictly the motor goal m . Starting at $t = 4$, the w^{FB} connections now carry a memory ($t - 1$) signal about the motor activity in the output layer. The customized BPTT weight update equations are presented below:

$$\begin{aligned}
\delta_k^{FB}(t_{max}) &= \delta_k^{out}(t_{max}) \\
\delta_j^{hh}(t_{max}) &= f'(H_j(t_{max})) \sum_k w_{jk}^O \delta_k^{out}(t_{max}) \\
\delta_k^{FB}(t-1) &= \delta_k^{out}(t-1) + f'(O_k(t-1)) \left(\sum_j w_{jk}^{FB} \delta_j^{hh}(t) \right) \\
\delta_j^{hh}(t-1) &= \left(\sum_k w_{kj}^O \delta_k^{FB}(t-1) + \sum_j w_{jm}^{hh} \delta_m^{hh}(t) \right) f'(H_j(t-1)) \\
\Delta w_{jk}^{out} &= \sum_{t=t_{max}}^1 \delta_k^{out}(t) H_j(t), w_{jk}^{out} = w_{jk}^{out} + \eta \Delta w_{jk}^{out} \\
\Delta w_{jm}^{hh} &= \sum_{t=t_{max}}^2 \delta_j^{hh}(t) H_m(t-1), w_{jm}^{hh} = w_{jm}^{hh} + \eta \Delta w_{jm}^{hh} \\
\Delta w_{jk}^{FB} &= \sum_{t=t_{max}}^2 \delta_j^{hh}(t) O_k(t-1), w_{jk}^{FB} = w_{jk}^{FB} + \eta \Delta w_{jk}^{FB} \\
\Delta w_{ij}^R &= \delta_j^{hh}(t=3) R_i(t=3), w_{ij}^R = w_{ij}^R + \eta \Delta w_{ij}^R \\
\Delta w_j^C &= \sum_{t=t_{max}}^1 \delta_j^{hh}(t) C(t), w_j^C = w_j^C + \eta \Delta w_j^C. \quad (6)
\end{aligned}$$

Maximum cross-correlation shift angle

A hidden unit is considered to have sensory-like behavior if its response with respect to the position of the visual cue does not change in anything but amplitude for two different contextual cues. In contrast, the hidden unit is more motor-like if the peak of its tuning curve (mapped with the position of visual stimulus on the x -axis) shifts between the two context values. To establish whether the hidden units developed more sensory (nonshifting response for all ω) or more motor (full or partial shifting for all ω) responses, we used cross-correlation (Eq. 7) in the following manner. The response of each hidden unit was mapped for ϕ ranging from 0 to 2π and for ω equal to (0, 45, 90, and 180°).

$$C(H_1, H_2) = \frac{\sum_{i=1}^n (H_{1i} - \bar{H}_1)(H_{2i} - \bar{H}_2)}{\sqrt{\sum_{i=1}^n (H_{1i} - \bar{H}_1)^2} \sqrt{\sum_{i=1}^n (H_{2i} - \bar{H}_2)^2}}. \quad (7)$$

Cross-correlation coefficients C were calculated for the tuning curves at $\omega_1 = 0^\circ$ (H_1) and $\omega_2 = 45, 90,$ and 180° (H_2) sampled at points $k \times 5^\circ$ ($k = 0, \dots, 72$). The index i in Equation 7 refers to the number of points that sample the tuning curve H (in our case, 73). The range of shifts was not limited, because the tuning curves were periodic. The maximum cross-correlation coefficient was chosen to denote the shift angle for which the two curves overlap the best. As a method, cross-correlation is only sensitive to alignment of tuning curves, and it does not depend on their exact functional form. An additional advantage is that it only indicates horizontal shifts and it does not depend on vertical shifts (gain changes). The cross-correlation coefficient was not calculated if one of the hidden units (either H_1 or H_2) had no response (or complete saturation) for the particular ω remapping. This was done by requiring that the firing rate H has $[\max(H) - \min(H)]/\max(H) > 0.1$ and $\max(H) > 0.2$. These criteria were established based on the empirical observations.

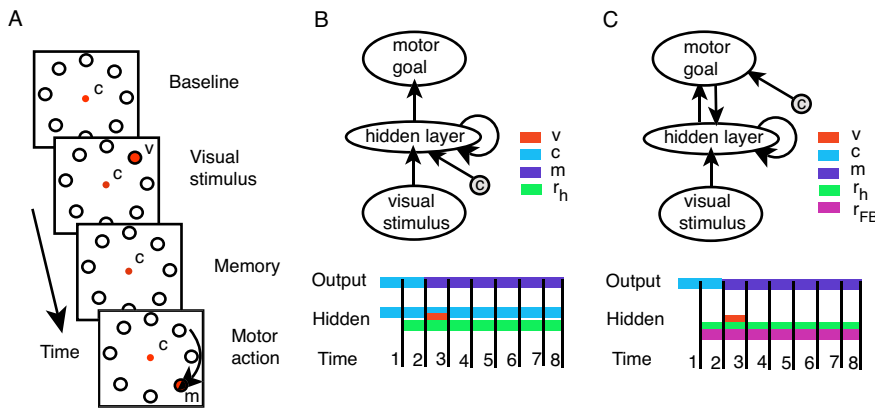


Figure 1. *A*, The timeline of the events in the trial. The trial starts with the display of a central fixation point in blue, green, red, or yellow, which corresponds to the rotation angle of 0, 45, 90, and 180°. This point is visible throughout the trial. In the second time period, the visual stimulus is briefly flashed in the periphery. The third frame corresponds to the so-called “memory” period when the visual stimulus is not visible anymore and the subject is waiting for the “go” signal to execute the motor action. The fourth frame represents the time period when the motor action is executed. This may be either a reach or a saccade toward the remapped position of the original visual stimulus. *B*, The schematic outline for network I. The input layer encodes the position of the visual stimulus. The hidden layer receives direct input from both the sensory layer as well as from the remapping rule unit *C*. The hidden layer units are recurrently connected. The output (motor) layer receives the hidden layer activation as the input. Color-coded time activities for the hidden and the output layers are shown in the bottom of the panel. Each color represents a parameter directly encoded in the hidden and the output layers: *c* is the remapping rule, *v* is the visual stimulus location, *m* is the remapped position of the stimulus (motor goal), and r_h is the recurrent activity in the hidden layer. *C*, The schematic outline for network II. The only difference from network I is that the rule information *C* now enters directly into the output layer. The notion about the rule comes to the hidden layer via the feedback connection from the output layer. The timeline of the network activity is given in the bottom of network II schematic. Note that in the first time step, there is no input coming into the hidden layer. Meanwhile, the rule information enters the output layer. From the second time step, the feedback connections r_{FB} bring in the context information into the hidden layer. Also, the hidden layer units start receiving the memory activity about their previous activation via r_h .

Extended network design

Hybrid network III. A third network combines the architectures of the previous two models. It is identical to network I, except it also has a feedback from the output layer as network II. The idea is to understand what pure motor feedback does to the behavior of the hidden units when the context arrives via separate (potentially sensory) pathways. The timeline of the network training and the methods are the same as previously explained.

Results

Both network architectures (Fig. 1*B,C*) successfully converge to properly encode the context-specific visuomotor mapping after learning. Figure 2 displays response maps of the unimodal activity in all eight output units in network I. These responses represent the “late memory” period ($t = 8$), and the activities are mapped with respect to the position of the visual cue (*x*-axis: -180 to $+180^\circ$; 5° increments) and four remapping conditions (*y*-axis: rule value of 0.25, 0.5, 0.75, and 1 corresponding to 90, 0, 180, and 45° clockwise rotation, respectively). A particular output unit is active every time the remapped position of the visual target falls within the preferred 1D tuning direction of the unit. For example, the fifth output unit has its preferred direction at 0° . This means that when the remapping rule dictates 90° clockwise rotation, only visual targets that are positioned in the vicinity of $+90^\circ$ will excite this neuron. Similar reasoning applies for all other rotation angles, and it explains the shifting pattern. The output pattern for network II looks identical, because the training goal was the same for both networks.

Hidden layer units acquired their tuning as a result of the network training. In contrast to the output layer, the resulting tuning was not directly constrained by the teaching signal. As a

consequence, the two network architectures converged onto different population encodings in the hidden layer, while both functionally achieved the same spatial visuomotor mapping between sensory input and motor output. Figure 3*A* shows activation in five representative example units of the hidden layer of network I at the end of the memory period ($t = 8$). Overall, a majority of units has a distinct unimodal tuning with respect to the position of the visual stimulus (*x*-axis), where the center of the tuning function is independent of the remapping rule (*y*-axis). Different remapping conditions result in the increase or decrease in the activation of the unit, additionally followed by the widening/narrowing of the tuning curve. There is very little shifting in the ϕ preferred direction, in contrast to the output layer. We therefore call the predominant tuning properties in the hidden layer of network I “context modulated visual memory tuning.” In contrast, preferred directions of hidden units in network II as a function of cue position strongly depend on the remapping rule (Fig. 3*B*). In this respect, the hidden units are similar to the motor-tuned output layer. Additionally, responses of hidden units in network II also show some context modulation.

To quantify the predominance of visual versus motor tuning, we analyzed the extent of the shift in preferred directions of all hidden units using cross-correlation coefficients (see Materials and Methods). Figure 3, *C* and *D*, shows the distribution of angular shifts, which correspond to the maximum correlation coefficient for hidden units in network I/II. The red histogram in Figure 3 represents alignment angles between the mapping conditions $\omega = 0^\circ$ and $\omega = 45^\circ$. The blue histogram is the equivalent measure for $\omega = 0^\circ$ and $\omega = 90^\circ$, whereas the gray histogram shows the shifts between the $\omega = 0^\circ$ and $\omega = 180^\circ$ conditions. The means (and SDs) of all three distributions ($-4.1 \pm 29.0^\circ$, $-0.3 \pm 42.2^\circ$, and $0.9 \pm 21.3^\circ$, respectively) suggest that, for network I, very little shifting of the tuning curves as a function of ω occurs (i.e., the tuning is independent of the context) (Fig. 3*C*). The results are different for network II (Fig. 3*D*). In the case of the pure motor behavior of the hidden units, they would have shifted the peak of their tuning curves by the value of ω (i.e., $45, 90,$ and 180°). Distribution means ($-42.1 \pm 21.9^\circ$, $-60.0 \pm 51.1^\circ$, and $-173.6 \pm 73.8^\circ$) in fact suggest motor tuning, although the deviation from the expected values and the large SDs indicate that the shifting is only partial.

The difference in the hidden layer population encoding of networks I and II can plausibly be explained by the network architecture. The motor tuning in the hidden layer of network II is a direct consequence of the feedback projections from the output layer (w^{FB}), which were shaped during the network training. Each hidden unit became “driven” by a systematic subset of output units, namely those output units that have corresponding tuning. This relationship between output and hidden units could be quantified by the same cross-correlation analysis as already used in Figure 3, *C* and *D*. This time, the cross-correlation index is calculated for the shifting overlap between the hidden

and output unit for $\omega = 0^\circ$ condition. The feedback weights (w^{FB}) are plotted (Fig. 4) with respect to the shift angle $\Delta\Phi$, which maximizes the cross-correlation between the hidden and output unit. When the hidden units are mainly driven via feedback, then the w^{FB} weight should be positive in case the hidden unit has a similar preferred direction and negative in case the hidden unit has an opposite preferred direction ($\Delta\Phi \approx \pm 180^\circ$).

It is possible that some network II units acquire the same visual memory behavior as network I. This happens when the w^{FB} weights, which connect the eight output units with a hidden unit, remain untuned after the network training. This particular hidden unit then always receives relatively small

$$\sum_l w_{jl}^{FB} O_l(t-1)$$

signal from the output layer. The only other “driving” signal is the recurrent input from the previous time step, which is predominantly shaped by the visual cue memory. In a rare case (1 of 10 network trainings with random initializations of weights), a network II developed predominantly visual memory.

The design of our networks allows a simplistic analysis of the dynamic encoding of visuomotor transformations by analyzing the time course of tuning in the hidden layer. In our studies, the remapping rule (context) information was always presented first and followed by a transient visual stimulus presentation. Figure 5A shows the activity of one exemplary output unit across all eight time steps of the activity in network I (equivalent in network II). The output units were trained to encode the value of the remapping rule at the first two time steps. As soon as the visual cue information becomes available, the output units are required to represent the motor goal (from $t \geq 3$). Figure 5B shows an example hidden unit in network I. Although the spatial tuning of the hidden units is quite different from the output units, the dynamics of tuning are similar. All hidden units encode only the rule signal at $t = 1$ and $t = 2$, and they acquire their characteristic spatial tuning curve as soon as the visual cue is presented at $t = 3$. The sensory-context integration only happens at $t = 3$. Everything after is a “memory” for this computation, which is kept in the network via the recurrent connections in the hidden layer. Note that the hidden units keep receiving a direct contextual input throughout the trial (Eq. 3). This context signal is obviously very weak compared with the recurrent




Figure 2. Motor tuning in the output layer. Activation in eight Gaussian units (1D) with respect to the position of visual cue (x) and the remapping rule value (y). The activation is expressed in arbitrary units, and it ranges from 0 to 1.




Figure 3. *A*, Activation of five example units in the hidden layer of network I. There are 40 hidden units, each mapped for the full range of the input stimuli values (from -180 to $+180^\circ$) in x and four rule values y (0.25, 0.5, 0.75, and 1), which correspond to the rotation angles of 90 , 0 , 180 , and 45° . *B*, Activation of five example units in the hidden layer of network II, which contains 40 hidden units. *C*, Maximum correlation angle shift (alignment angle) for the hidden units in network I. *D*, Maximum correlation angle shift for the hidden units in network II.

$$\sum_m w_{mj}^H H_m(t-1)$$

input, because it is visible only through a very small modulation (in the width and the amplitude) of the tuning curve as the time progresses.

Figure 5C shows the corresponding tuning dynamics for hidden units in network II. At $t = 1$, the hidden units receive zero




Figure 4. The weight matrix w^{FB} (8×40) in network II plotted as a function of the difference between the preferred direction of the hidden and output units. Marked in red are the w^{FB} values for the five examples of motor-tuned hidden units in Figure 3B.

input (which results in uniform activation of 0.5 because of the sigmoid transfer function). At $t = 2$, the rule information enters the hidden layer indirectly, via feedback from the output layer:

$$\sum_l w_{jl}^{FB} O_l(t = 1).$$

Units also receive a nonspatial signal about the previous (uniform) activity of the hidden layer:

$$\sum_m w_{jm}^H H_m(t = 1).$$

As a consequence, the different remapping rules result in four different levels of activation in the hidden units. At $t = 3$, the visual cue information

$$\sum_i w_{ji}^R R_i(t = 3)$$

enters the hidden layer and is integrated with the context signal. Because of the nonlinear additive computation of the neuron, this resulted in context-modulated visual tuning. A memory about this visual tuning is present at $t = 4$ via w^H . However, this memorized visual representation gets direct competition from the top-down signal being back-projected from the output layer via w^{FB} . The resulting mixture shapes the activation of the hidden layer at $t = 4$. In network II, the feedback “motor memory” signal dominates the net input of the hidden units. This is a consequence of the specific magnitude and pattern of the w^{FB} weights as discussed in Figure 4.

In the extended study, a hybrid network was created (Fig. 6A). Network III was identical to network I regarding where the context entered the sensorimotor process. Additionally, a feedback from the output layer was included (as in network II). Hidden units developed predominantly motor-like behavior, as in network II. Figure 6B shows the same cross-correlation shift analysis as before. Distribution means for the ω of -45 , -90 , and -180°

correspond to $-37.7 \pm 40.8^\circ$, $-65.6 \pm 49.8^\circ$, and $-179.6 \pm 72.7^\circ$, suggesting motor-like tuning.

Discussion

We investigated mechanisms underlying rule-based sensorimotor transformations, which include tasks such as contextually guided saccades and reaches. Two networks with different architectures were trained to simulate context-specific visuomotor transformations. One extrinsic topographical map (visual input) had to be mapped onto another (motor-goal in visual coordinates). The spatial mapping-rule depended on a nonspatial, contextual input signal. Both networks effectively implemented a gain-modulatory mechanism for context integration but converged onto different predominant encoding schemes in the hidden layer.

We consider network II to more likely represent cortical sensorimotor processing. First, the connectivity pattern in combination with the potential sources of context information in the cortex makes network II more plausible on anatomical-physiological grounds. Second, the predominant motor-goal representations in the hidden layer of network II, compared with network I, resemble the encoding in posterior parietal cortex (PPC) during the memory period of visuomotor tasks similar to the one simulated here. Third, network II reproduces qualitatively the dynamics of context-specific sensorimotor transformations as observed in cortical sensorimotor areas. However, based on our hybrid network study (Fig. 6), we cannot completely rule out that the context does not arrive into the PPC via feedforward connections. We discuss these arguments below.

The mechanism we propose for sensory-context integration is very general. It could apply to multiple sensorimotor modalities. To keep our discussion compact, especially regarding anatomical-physiological interpretations, we focus on visuomotor transformations for reaching only. Also, in our models, we do not include transformations from extrinsic motor-goal representations to intrinsic motor commands. Detail studies on how cortical circuitry may implement kinematically correct motor commands for reaches and saccades can be found in other studies (Crawford and Guitton, 1997; Crawford et al., 2004; Smith and Crawford, 2005).

Anatomical-physiological interpretations of the network connectivity

It is yet unclear where in the brain and how contextual information is integrated with spatial sensory information to achieve rule-guided sensorimotor processing (e.g., to allow flexible movement behavior in a given sensory environment). Areas typically associated with reach planning are parts of the PPC, like the parietal reach region (PRR) or area 5, and dorsal premotor cortex (PMd) (Mountcastle et al., 1975; Snyder et al., 1997; Caminiti et al., 1998; Batista et al., 1999; Battaglia-Mayer et al., 2000; Buneo et al., 2002). Extensive experimental evidence suggests that areas such as the PFC (Petrides, 1982; Dias et al., 1996; White and Wise, 1999; Miller and Cohen, 2001; Wallis et al., 2001) provide behavioral context information to control motor-goal selection and movement execution. The underlying idea is that the premotor cortex integrates sensorimotor information, mediated via frontoparietal loops, with contextual information, mediated via prefrontal-premotor networks, as reflected in network II. However, context (e.g., mapping rule) signals have also been reported recently in PPC, and motor-goal representations suggest context-specific sensorimotor processing in this area (Stoet and Snyder, 2004; Gail and Andersen, 2006). In principle, this could indicate

more local context integration by direct combination of spatial and contextual sensory cues within PPC, as reflected in network I and the hybrid network from Figure 6A. We argue that context selectivity and motor-goal tuning in PPC do not contradict the former interpretation of a sensory-context integration in frontal areas. Rather, they are likely the consequence of strong top-down signals from motor-tuned structures (see below).

Our two basic network models, as well as our intermediate model, simulate the alternative views on sensory-context integration in a very simplified way. We investigated the spatiotemporal behavior of the networks when the contextual information is either fed into the hidden layer in a direct feedforward manner like sensory input (Figs. 1B, 6A) or, alternatively, when the context is mediated via top-down feedback from the output layer (Fig. 1C). We assumed that the hidden layer corresponds to the posterior parietal cortex (Zipser and Andersen, 1988; Pouget and Sejnowski, 1995; Xing and Andersen, 2000a; Smith and Crawford, 2005), and that the output layer represents motor-tuned subpopulations of the PMd. It is also possible that the output layer resides within PRR, only representing a higher stage of processing within the area. Finally, other parietal areas with reach-related activity could represent this output stage (e.g., area 5, which appears to be downstream of PRR, at least in terms of coordinate frame representation of reach targets) (Buneo et al., 2002; Buneo and Andersen, 2006).

Strong frontoparietal projections from premotor to posterior parietal cortex, like the top-down projections in network II and the hybrid network, are well established (Pandya and Yeterian, 1990; Barbas and Pandya, 1991; Goldman-Rakic, 1998; Petrides and Pandya, 2002). Note that, in network II, the feedback projection carries both a contextual signal as well as a motor memory signal, whereas in the hybrid network, these connections only carry a motor memory signal.

Prefrontal projections directly to the posterior superior parietal lobe are rather weak, if they at all exist to a reasonable extent (Petrides and Pandya, 2002). The latter would be required to explain the direct projection of context information (presumably from PFC) to the hidden layer (presumably PPC) in network I and hybrid network. Alternatively, in highly trained subjects, the direct context input to PPC might be mediated independent from PFC via sensory areas (Grol et al., 2006). This processing scheme would also be compatible with the architecture of either network I or the hybrid network (the source of the context signal being sensory rather than prefrontal areas). But purely sensory representations of nonspatial context information (e.g., indicated by a color cue) have not been found for PPC (Stoet and Snyder, 2004; Gail and Andersen, 2006), making this alternative unlikely. Together, this makes the network II architecture more plausible.

Context integration by gain modulation

At the time of the spatial cue presentation (with the context cue already present), both of our networks implement integration of the sensory $R(\phi)$ and contextual $C(\omega)$ inputs via an approximate




Figure 5. Temporal profile for network I/II, $t = 1, \dots, 8$. **A**, Activity in time for the fifth output unit. **B**, Activity in time for one example hidden unit from Figure 3A. **C**, The activity in time for one example hidden unit from Figure 3B.

gain mechanism (Eq. 3). As long as the net input remains within the linear part of the sigmoidal transfer function f , the response only changes in amplitude (pure multiplicative gain). Once the net input encounters the upper nonlinear part (and further saturation region) of the sigmoid, the response also widens. The opposite (nonlinear narrowing) happens on the other end of the sigmoidal curve.

Sensory-context integration via gain modulation has been proposed previously in a basis network model (Salinas, 2004). The basis units had gain modulation explicitly built-in by a multiplicative term that could take a discrete set of preassigned context values. Otherwise, the input units encoded a discrete set of 16 “parametric visual” stimuli in a way that did not represent a spatial topographic map. Our models instead solve the task of flexibly remapping one topographic representation onto another. Because the input and output units in our networks had overlapping tuning curves, they define neighborhood relationships (topology) and are suitable to encode retinotopic stimulus/motor-goal position. The overall conclusion in the study by Salinas (2004) and in our study is that the context modulates the sensory tuning curves in a similar way that the classic “gain” factors such as gaze direction and head position do (Andersen et al., 1990; Brotchie et al., 1995; Cohen and Andersen, 2000).

Motor-goal tuning in parietal sensorimotor areas

In our models, successful sensory-context integration did not depend on whether the context information arrived directly or indirectly into the hidden layer. However, the predominant motor-goal tuning in the hidden layer was closely tied to the information flow within the network and occurred in both network II and the hybrid network during the memory period. This may provide a theoretical basis for interpreting the recent experimental findings (Gail and Andersen, 2006), which reported that the contextual signal induced motor-goal tuning as opposed to

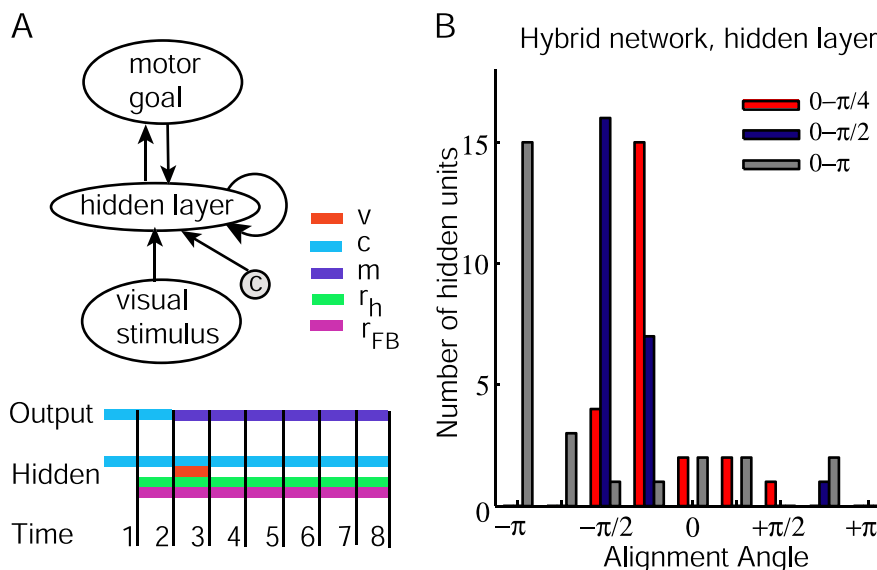


Figure 6. *A*, Schematics of the activity for the hybrid network III. *B*, Maximum correlation angle shift (alignment angle) for the hidden units in the hybrid network III.

gain modulated sensory tuning in the PPC during memory-guided remapping tasks. Gail and Andersen (2006) also reported on the context information being present in the PPC even before the presentation of the visual stimulus. This condition was satisfied in all of our networks, if we choose to interpret the hidden network layer as a simplified equivalent of the PPC.

We propose that the motor-like behavior of the neurons in the PPC might reflect the fact that these neurons receive a strong feedback signal from motor-encoding structures in the frontal lobe, or even within the parietal lobe.

Dynamics of context-specific sensorimotor transformations

The recurrent connections in our networks allowed the simulation of a temporal aspect of sensory-context integration. We do not account for any physiological time constants, and thus it is not possible to quantitatively compare the dynamics of the artificial units with physiological neuronal or behavioral responses. Also, the information propagates instantly from input to output layers in each time step: as soon as the cue information is available to the input units, the output unit activity is also updated. Nevertheless, qualitative observations can be made.

A recent experimental study (Gail and Andersen, 2006) investigated how context is combined with sensory information in the PRR, a brain area associated with sensory integration based on the gain mechanism (Batista et al., 1999; Buneo et al., 2002). This study focused on PRR activity in the transition phase from visual cue presentation to the “memory period,” a time segment in which a monkey was preparing to execute contextually guided pro- or anti-reach with respect to the memorized cue position. Results showed that visuomotor transformations are performed in a context-specific manner resulting in a predominant motor-goal representation during the memory period in PRR neurons. Although this study focuses on sensorimotor transformations for visually guided reaching, the proposed neuronal mechanism is general and may apply to other modalities, such as visually guided saccades (Schlag-Rey et al., 1997; Amador et al., 1998; Gottlieb and Goldberg, 1999; Zhang and Barash, 2000; Amador et al., 2004).

Hidden units change their tuning while in the transition from

cue to memory period in network II (Fig. 5C). Activity that previously encoded only spatial cue (and context) now also reflects the motor-goal representation. The network was trained to treat all four context transformations equally, and thus the activation timeline for the congruent ($\omega = 0^\circ$) mapping is the same as for the incongruent ($\omega \neq 0^\circ$) mapping cases. However, it is possible that a biological neuron would take some time to switch its default congruent tuning, inherited from the cue period, into incongruent tuning. Such a latency difference in motor-goal tuning was found in PRR (Stoet and Snyder, 2004; Gail and Andersen, 2006). It corresponds to a large number of behavioral findings indicating slower reaction times whenever a stimulus-response mapping is spatially incongruent or, more generally, when stimulus and response share a feature (here, space) but with low compatibility (for review, see Kornblum et al., 1990; Proctor and Vu, 2002). A recent

computational study on target selection (Cisek, 2006) suggests that a decision influenced by a sensory cue first appears in PPC and then further propagates to PMd. In contrast, a decision based on abstract rules possibly first appears in the frontal regions from where it propagates to PPC. This view is consistent with our network II, where the motor-goal (decision) information is inherited by PPC from the output stage. Different from dynamic field models for movement preparation (Erlhagen and Schoner, 2002; Cisek, 2006), our models developed their connectivity patterns by learning and are able to produce motor goals at locations that were not previously cued by a sensory stimulus (anti-reach).

Conclusion

Based on neural network simulation, we propose that the integration of sensory and contextual cues in parietal cortex happens via a gain-modulation mechanism. The motor-like behavior of units in the parietal cortex after sensory-context integration could be explained by the existence of the strong feedback connections from motor output stages. It is possible that this feedback initially carries the context signal to the parietal cortex after originating from prefrontal areas and being mediated by premotor areas. This would help explain the presence of a high-level signal in posterior parietal cortex without a direct projection from prefrontal cortex, a likely candidate for the source of abstract context information.

References

- Amador N, Schlag-Rey M, Schlag J (1998) Primate antisaccades. I. Behavioral characteristics. *J Neurophysiol* 80:1775–1786.
- Amador N, Schlag-Rey M, Schlag J (2004) Primate antisaccade. II. Supplementary eye field neuronal activity predicts correct performance. *J Neurophysiol* 91:1672–1689.
- Andersen RA, Buneo CA (2003) Sensorimotor integration in posterior parietal cortex. In: *The parietal lobes* (Siegel AM, Andersen RA, Freund H-J, Spencer DD, eds), pp 159–177. Philadelphia: Lippincott Williams and Wilkins.
- Andersen RA, Bracewell RM, Barash S, Gnadt JW, Fogassi L (1990) Eye position effects on visual, memory, and saccade-related activity in areas LIP and 7a of macaque. *J Neurosci* 10:1176–1196.
- Barbas H, Pandya DN (1991) Patterns of connections of the prefrontal cortex in the rhesus monkey associated with cortical architecture. In: *Frontal*

- lobe function and dysfunction (Levin HS, Eisenberg HM, Benton AL, eds), pp 35–58. Oxford: Oxford UP.
- Batista AP, Buneo CA, Snyder LH, Andersen RA (1999) Reach plans in eye-centered coordinates. *Science* 285:257–260.
- Battaglia-Mayer A, Ferraina S, Mitsuda T, Marconi B, Genovesio A, Onorati P, Lacquaniti F, Caminiti R (2000) Early coding of reaching in the parietooccipital cortex. *J Neurophysiol* 83:2374–2391.
- Brotchie PR, Andersen RA, Snyder LH, Goodman SJ (1995) Head position signals used by parietal neurons to encode locations of visual stimuli. *Nature* 375:232–235.
- Buneo CA, Andersen RA (2006) The posterior parietal cortex: sensorimotor interface for the planning and online control of visually guided movements. *Neuropsychologia* 44:2594–2606.
- Buneo CA, Jarvis MR, Batista AP, Andersen RA (2002) Direct visuomotor transformations for reaching. *Nature* 416:632–636.
- Caminiti R, Ferraina S, Mayer AB (1998) Visuomotor transformations: early cortical mechanisms of reaching. *Curr Opin Neurobiol* 8:753–761.
- Cisek P (2006) Integrated neural processes for defining potential actions and deciding between them: a computational model. *J Neurosci* 26:9761–9770.
- Cohen YE, Andersen RA (2000) Eye position modulates reach activity to sounds. *Neuron* 27:647–652.
- Crawford JD, Guitton D (1997) Visual-motor transformations required for accurate and kinematically correct saccades. *J Neurophysiol* 78:2811–2816.
- Crawford JD, Medendorp WP, Marotta JJ (2004) Spatial transformations for eye–hand coordination. *J Neurophysiol* 92:10–19.
- Dias R, Robbins TW, Roberts AC (1996) Dissociation in prefrontal cortex of affective and attentional shifts. *Nature* 380:69–72.
- Erlhagen W, Schoner G (2002) Dynamic field theory of movement preparation. *Psychol Rev* 109:545–572.
- Gail A, Andersen RA (2006) Neural dynamics of context-specific sensorimotor transformations in monkey parietal reach region. *J Neurosci* 26:9376–9384.
- Goldman-Rakic PS (1998) Topography of cognition-parallel distributed networks in primate association cortex. *Annu Rev Neurosci* 11:137–156.
- Gottlieb J, Goldberg ME (1999) Activity of neurons in the lateral intraparietal area of the monkey during an antisaccade task. *Nat Neurosci* 2:906–912.
- Grol MJ, de Lange FP, Verstraten FA, Passingham RE, Toni I (2006) Cerebral changes during performance of overlearning arbitrary visuomotor associations. *J Neurosci* 26:117–125.
- Jaeger H (2002) A tutorial on training recurrent neural networks, covering BPTT, RTRL, EKF and the “echo state network” approach. In: GMD report 159: German national research center for information technology.
- Kalaska JF (1996) Parietal cortex area 5 and visuomotor behavior. *Can J Physiol Pharmacol* 74:483–498.
- Kornblum S, Hasbroucq T, Osman A (1990) Dimensional overlap: cognitive basis for stimulus-response compatibility—a model and taxonomy. *Psychol Rev* 97:253–270.
- Miller EK, Cohen JD (2001) An integrative theory of prefrontal cortex function. *Annu Rev Neurosci* 24:167–202.
- Mountcastle VB, Lynch JC, Georgopoulos A, Sakata H, Acuna C (1975) Posterior parietal association cortex of the monkey: command functions for operations within extrapersonal space. *J Neurophysiol* 38:871–908.
- Pandya DN, Yeterian EH (1990) Prefrontal cortex in relation to other cortical areas in rhesus monkey: architecture and connections. *Prog Brain Res* 85:63–94.
- Pasupathy A, Miller EK (2005) Different time courses for learning-related activity in the prefrontal cortex and striatum. *Nature* 433:873–876.
- Petrides M (1982) Conditional associative-learning after selective prefrontal lesions in the monkey. *Behav Brain Res* 5:407–413.
- Petrides M, Pandya DN (2002) Association pathways of the prefrontal cortex and functional observations. In: *Principles of frontal lobe function* (Stuss DT, Knight RT, eds), p 31. Oxford: Oxford UP.
- Platt ML, Glimcher PW (1999) Neural correlates of decision variables in parietal cortex. *Nature* 400:233–238.
- Pouget A, Sejnowski TJ (1995) Spatial representations in the parietal cortex may use basis functions. In: *Advances in neural information processing systems*, pp 157–164. Cambridge, MA: MIT.
- Proctor RW, Vu KP (2002) Mixing location-irrelevant and location-relevant trials: influence of stimulus mode on spatial compatibility effects. *Mem Cognit* 30:281–293.
- Salinas E (2004) Fast remapping of sensory stimuli onto motor actions on the basis of contextual modulation. *J Neurosci* 24:1113–1118.
- Salinas E, Abbott LF (1996) A model of multiplicative neural responses in parietal cortex. *Proc Natl Acad Sci USA* 93:11956–11961.
- Salinas E, Thier P (2000) Gain modulation: a major computational principle of the central nervous system. *Neuron* 27:15–21.
- Schlag-Rey M, Amador N, Sanchez H, Schlag J (1997) Antisaccade performance predicted by neuronal activity in the supplementary eye field. *Nature* 390:398–401.
- Shenoy KV, Bradley DC, Andersen RA (1999) The influence of gaze rotation upon the visual response of primate MSTd neurons. *J Neurophysiol* 81:2764–2786.
- Smith M, Crawford J (2005) Distributed population mechanism for the 3-D oculomotor reference frame transformation. *J Neurophysiol* 93:1742–1761.
- Snyder LH, Batista AP, Andersen RA (1997) Coding of intention in the posterior parietal cortex. *Nature* 386:167–170.
- Stoet G, Snyder LH (2004) Single neurons in posterior parietal cortex of monkeys encode cognitive set. *Neuron* 42:1003–1012.
- Takeda K, Funahashi S (2002) Prefrontal task-related activity representing visual cue location or saccade direction in spatial working memory task. *J Neurophysiol* 87:567–588.
- Wallis JD, Miller EK (2003) From rule to response: neuronal processes in the premotor and prefrontal cortex. *J Neurophysiol* 90:1790–1806.
- Wallis JD, Anderson KC, Miller EK (2001) Single neurons in prefrontal cortex encode abstract rules. *Nature* 411:953–956.
- Werbos PJ (1990) Backpropagation through time: what it does and how to do it. *Proc IEEE* 78:1550–1560.
- White IM, Wise SP (1999) Rule-dependent neuronal activity in the prefrontal cortex. *Exp Brain Res* 126:315–335.
- Wise SP, Boussaoud D, Johnson PB, Caminiti R (1997) Premotor and parietal cortex: corticocortical connectivity and combinatorial computations. *Annu Rev Neurosci* 20:25–42.
- Xing J, Andersen RA (2000a) Models of the posterior parietal cortex which perform multimodal integration and represent space in several coordinate frames. *J Cogn Neurosci* 12:601–614.
- Xing J, Andersen RA (2000b) The memory activity of LIP neurons in sequential eye movements simulated with neural networks. *J Neurophysiol* 84:651–665.
- Zhang M, Barash S (2000) Neuronal switching of sensorimotor transformations for antisaccades. *Nature* 408:971–975.
- Zipser D, Andersen RA (1988) A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature* 331:679–684.