

Paying Attention through Eye Movements: A Computational Investigation of the Premotor Theory of Spatial Attention

Marco Casarotti*, Matteo Lisi*, Carlo Umiltà, and Marco Zorzi

Abstract

■ Growing evidence indicates that planning eye movements and orienting visuospatial attention share overlapping brain mechanisms. A tight link between endogenous attention and eye movements is maintained by the premotor theory, in contrast to other accounts that postulate the existence of specific attention mechanisms that modulate the activity of information processing systems. The strong assumption of equivalence between attention and eye movements, however, is challenged by demonstrations that human observers are able to keep attention on a specific location while moving the eyes elsewhere. Here we investigate whether a recurrent model of saccadic planning can account for attentional effects without requiring

additional or specific mechanisms separate from the circuits that perform sensorimotor transformations for eye movements. The model builds on the basis function approach and includes a circuit that performs spatial remapping using an “internal forward model” of how visual inputs are modified as a result of saccadic movements. Simulations show that the latter circuit is crucial to account for dissociations between attention and eye movements that may be invoked to disprove the premotor theory. The model provides new insights into how spatial remapping may be implemented in parietal cortex and offers a computational framework for recent proposals that link visual stability with remapping of attention pointers. ■

INTRODUCTION

The premotor theory of spatial attention (Rizzolatti, Riggio, & Sheliga, 1994; Umiltà, Riggio, Dascola, & Rizzolatti, 1991; Rizzolatti, Riggio, Dascola, & Umiltà, 1987) maintains that endogenous (i.e., top-down) orienting of visuospatial attention originates from the activation of the cortical circuits involved in saccadic planning. Preparation of a saccadic movement produces, by means of recurrent projections from premotor areas to parietal spatial maps, a processing facilitation for stimuli located in the region of space toward which the motor program is prepared. Planning a saccade is equivalent to shifting attention in space, because a covert movement of attention occurs when an eye movement is prepared but not executed. In contrast, other theories of spatial attention postulate the existence of specific attention mechanisms that modulate the activity of information processing systems (see, e.g., Mesulam, 1990).

Neurophysiological data strongly support the premotor theory indicating that spatial attention is related to eye-movement planning structures, including the FEFs (Moore & Fallah, 2001, 2004; Moore, Armstrong, & Fallah, 2003) and the superior colliculus (SC; Muller, Philastides, &

Newsome, 2005; Ignashchenkova, Dicke, Haarmeier, & Thier, 2004; Kustov & Robinson, 1996). Reversible neuro-disruption of FEF, both in monkeys (Vardak, Ibos, Duhamel, & Olivier, 2006; Moore & Fallah, 2004) and humans (see Chambers & Mattingley, 2005, for a review), affects the orienting of spatial attention. Neurons in the intraparietal sulcus (IPs) generate action-oriented representations of space and are also crucially involved in the top-down (endogenous) control of spatial attention (see Colby & Goldberg, 1999, for a review). Neural activity in the lateral intraparietal area (LIP) depends on the spatial and temporal dynamics of attention (Bisley & Goldberg, 2003) and represents only salient targets, suggesting that LIP neurons generate a saliency map of the visual environment.

Neuroimaging studies indicate that top-down control of spatial attention in humans recruits a network of cortical areas including the IPs and the FEF (see Corbetta & Shulman, 2002, for a review). That is, the network of brain regions involved in endogenous orienting of spatial attention largely overlaps with the network subserving sensorimotor transformations for saccadic movements (Beauchamp, Petit, Ellmore, Ingeholm, & Haxby, 2001; Nobre, Gitelman, Dias, & Mesulam, 2000; Corbetta et al., 1998).

Recent behavioral data, however, challenge the premotor theory by showing dissociations between attention and eye movements. Golomb, Chun, and Mazer (2008) directly addressed the issue of how the topography of visuospatial

University of Padova

*These authors contributed equally and are listed in alphabetical order.

attention reorganizes after an eye movement. They developed a gaze-contingent paradigm in which participants performed an eye movement while keeping in memory the location of a spatial cue. Maintaining a location in memory, indeed, amounts to voluntarily deploy spatial attention to the memorized location (see Awh & Jonides, 2001, for a review). Results demonstrated that attention can be maintained on the location of a spatial cue while moving the eyes elsewhere. This should not be possible if control of eye movements and control of attention were tightly coupled. More specifically, the study revealed facilitation effects at both retinotopic and spatiotopic coordinates of the attended location around the time of an intervening saccade. Retinotopic facilitation prevailed for 100–200 msec after the eye movement, although this location was task-irrelevant. Conversely, at later delays, the attentional benefit prevailed at the spatial, task-relevant coordinates of the attended location.

These findings were replicated under different experimental manipulations (Golomb, Marino, Chun, & Mazer, 2011; Golomb, Pulido, Albrecht, Chun, & Mazer, 2010) and corroborated by neuroimaging evidence (Golomb, Nguyen-Phuc, Mazer, McCarthy, & Chun, 2010). According to Golomb et al. (2008), these results imply that the basic coordinate system of spatial attention is retinotopic and it must be updated to compensate for intervening eye movements. However, the data are also consistent with the alternative hypothesis that spatial attention operates on two saliency maps (one retinotopic and the other spatiotopic) with different time courses (see also Astle, 2009).

Spatial updating of attended locations is consistent with single-cell studies showing that LIP neurons update the representation of visual space across eye movements (Duhamel, Colby, & Goldberg, 1992). LIP neurons have retinotopic receptive fields (RFs) and carry visual and visual memory signals. Spatial representations in LIP, however, are not simply retinotopic. Indeed, remembered target locations are remapped in the coordinates of the new fixation point after an eye movement. Some LIP neurons, moreover, anticipate the retinal consequences of intended eye movements by becoming transiently responsive to stimuli presented in their postsaccadic RF (i.e., predictive remapping).

Remapping in LIP updates the internal representation of visual space in conjunction with eye movements. This process requires a mechanism that produces a shift of activity from the original coordinate frame to the post-saccadic frame using oculomotor information. A corollary discharge (CD) of the saccadic command is supposed to originate in the SC, from which it gets to the FEF via the mediodorsal thalamus (see Sommer & Wurtz, 2008, for a review). FEF neurons in turn are functionally coupled with LIP (Ferraina, Pare, & Wurtz, 2002). CD signals may also reach LIP neurons without crossing the FEF via the lateral pulvinar nucleus (Clower, West, Lynch, & Strick, 2001). This distributed network is supposed to

fulfill the computation of vector subtraction, which permits to achieve spatial remapping without requiring an explicit suparetinal representation of target location. However, how the brain performs this computation remains unknown.

In the seminal article by Duhamel et al. (1992), remapping was attributed to shifting RFs. This account implies that each LIP cell should be connected to all locations on the retina through interneurons. During fixation, only the retinal location that corresponds to the classic RF can be accessed, whereas all the other locations are gated. Around the time of an eye movement, all RFs shift from their default location to the appropriate offset location, which depends on the current saccade target. The shifting RF model has been recently challenged on the basis of two compelling arguments (see Cavanagh, Hunt, Afraz, & Rolfs, 2010, for discussion). The first one takes advantage of cross-modal anticipatory responses, which are analogous to predictive remapping. In this case, no shifting RFs can be invoked, because rewiring should take place between different modalities. Second, the updating of remembered spatial locations in LIP rules out the hypothesis of shifting RFs, because at the time of remapping there is no activity on the retina or in earlier visual cortices. Cavanagh et al. (2010) argue that the only source for remapping must be a transfer of information from currently active cells that hold spatial locations in memory. This mechanism requires that horizontal connections can transfer activation across LIP cells using a corollary signal of the upcoming saccade.

If remapping involves activation transfer across a saliency map, one important question is what kind of connectivity might be involved. Quaia, Optican, and Goldberg (1998) proposed a computational model of LIP–FEF interactions that performed spatial remapping through horizontal connections in LIP. However, the model required specific connectivity and operations at the dendritic level, which are difficult to implement in a biological circuit. Horizontal connections were used also by Xing and Andersen (2000a) to model spatial updating in LIP. The connection weights, however, were computed using an optimization procedure with specific constraints. Moreover, the model included a set of memory units that stored one spatial location at a time. That is, it required as many memory buffers as targets to be stored. More recently, Keith and Crawford (2008) trained a back-propagation network to perform a double saccade task. After learning, the network achieved spatial remapping by means of a lateral displacement in the hidden units' RFs. However, back-propagation is not considered biologically plausible, because learning employs signals that are nonlocally available. Moreover, the model has a feed-forward architecture, whereas bidirectional propagation is a critical computational principle in the cerebral cortex (O'Reilly, 1998), where recurrent connections are ubiquitous.

Unlike back-propagation models, basis function (BF) networks with recurrent connectivity can be readily mapped onto parietal circuits (Pouget & Snyder, 2000, for a review). Indeed, the properties of posterior parietal neurons that combine sensory and posture signals suggest that they may serve as BFs with which the brain computes coordinate transformations. BFs are processing units that compute the product of nonlinear functions, which form their basis set, and a linear combination of their outputs is sufficient to approximate any arbitrary function of their inputs (Pouget & Sejnowski, 1997; Poggio, 1990). It follows that encoding space with BFs renders it possible to reduce nonlinear coordinate transformations to simple linear mappings. The resulting BF representation encodes spatial locations in a format that contains implicitly any frame of reference that can be derived from the input variables: for instance, a BF map that combines visual information with eye position contains a head-centered frame that can be read out with a simple linear transformation of the activity of the BF units (Pouget & Sejnowski, 1997). One drawback of the BF approach is the problem known as the curse of dimensionality: BF representations are subject to combinatorial explosion, because the number of units increases exponentially with the number of inputs being combined (for further discussion, see Pouget & Snyder, 2000). Nevertheless, the high redundancy of a BF representation can be exploited to optimally filter out noise in the sensory input (Deneve, Latham, & Pouget, 2001).

The BF approach is consistent with neurophysiological evidence showing that the activity of many parietal neurons involved in sensorimotor transformations approximates a multiplicative combination of sensory and posture signals (Andersen, 1989; Andersen, Essick, & Siegel, 1985). Cell encoding with multiplicative interaction of independent variables (i.e., gain-field coding) is considered as a major computational principle of nonlinear neuronal processing (Salinas & Thier, 2000, for a review). Computational studies determined how and under what conditions coordinate transformations can be performed by gain modulated neurons (Salinas & Abbott, 1995). How neurons combine their inputs in a directly multiplicative manner remains unclear, although a number of cellular mechanisms have been proposed (see Brozovic, Abbott, & Andersen, 2008). At the network level, gain modulation can arise as a consequence of learning rules that adjust the strength of synaptic connections to achieve specific coordinate transformations (Smith & Crawford, 2005; Xing & Andersen, 2000a, 2000b; Zipser & Andersen, 1988). Moreover, multiplicative responses can arise through population effects in a recurrent network with excitatory connections between similarly tuned neurons and inhibitory connections between differently tuned neurons (Salinas & Abbott, 1996). As a consequence, BFs can be seen as building blocks that simulate the activity of single gain

modulated neurons or population effects within many parietal cells.

Notably, recurrent BF networks are well suited for implementing internal forward models (Denève, Duhamel, & Pouget, 2007) that describe how sensory inputs are modified as a result of motor action. Growing empirical evidence suggests that the brain integrates sensory and motor signals using such internal models to perform a variety of tasks, such as predicting sensory information and optimal motor control (Todorov, 2004; Desmurget & Grafton, 2000; Kawato, 1999; Wolpert, Ghahramani, & Jordan, 1995). Because retinotopic representations change in a predictable way if the parameters of an eye movement are known, an internal forward model may be used for achieving spatial remapping across saccades (Vaziri, Diedrichsen, & Shadmehr, 2006).

This study aims to investigate whether a recurrent model of saccadic planning can account for attentional effects without requiring additional or specific mechanisms separate from the circuits that perform sensorimotor transformations for eye movements. Accordingly, attention orienting is implemented in terms of feedback effects because of saccadic planning and is explicitly concerned with action-oriented representations. The model builds on the BF approach and includes a circuit that achieves spatial remapping using an internal forward model of how visual signals are modified as a result of saccadic movements. The latter circuit provides new insight into how remapping operations may be implemented in parietal cortex and accounts for dissociations between attention and eye movements observed in gaze-contingent paradigms.

METHODS

Overview of the Model

In the spirit of a nested incremental modeling approach (Perry, Ziegler, & Zorzi, 2007), the model is built upon previous computational work on modeling sensorimotor transformations using BFs (Pouget & Snyder, 2000; Pouget & Sejnowski, 1997, for a review). The architecture of the model (Figure 1) consists of a BF map, which simulates the activity of LIP neurons, and a motor map that simulates saccadic planning in FEF through population coding. Each map has lateral connections that generate local excitation and long-range inhibition. This allows memory activity in the absence of visual input and competition between different population codes (Wang, 2001; Compte, Brunel, Goldman-Rakic, & Wang, 2000).

The BF map combines population codes representing retinal (r) and oculomotor (c) signals. As the neuron tuning curves are Gaussians centered at (r, c) , this layer is a two-dimensional radial BF map for retinal position and oculomotor command. Neurons are arranged topographically (e.g., Patel et al., 2010) along the corresponding




Figure 1. Recurrent neural network model for spatial attention. The model contains a radial BF map (LIP) for retinal position and oculomotor command and a motor map (FEF) encoding saccadic plans through population coding. Neurons in the LIP map combine visual and motor signals (i.e., retinal information and a CD of saccadic commands) to represent oculomotor space. Lateral connectivity in LIP implements an internal model that predicts the sensory consequences of saccadic movements. Each map has lateral connections with local excitation and long-range inhibition.

axis and are connected so as to estimate the remapped position of a memorized visual target across eye movements. As a result, given visual input r and oculomotor command c , the corresponding hill of activity in LIP will shift to the fixation neuron (i.e., coding for a 0° motor command) with preferred retinal position $r - c$. This recurrent connectivity implements an internal forward model that predicts the visual consequences of saccadic movements.

LIP neurons are reciprocally connected with FEF neurons through topographical projections. That is, an LIP neuron with preferred retinal position r is connected preferentially with an FEF neuron that codes for the corresponding target location. In agreement with the premotor theory, feedback of FEF activity to LIP neurons allows a motor program to generate endogenous, top-down attentional signals through the recruitment of neurons located upstream in parietal spatial maps. Moreover, the implementation of the circuit responsible for spatial remapping renders it possible to investigate the role of perisaccadic updating in attention orienting.

Recurrent Model: Implementation Details

The LIP map is composed of 441 neurons that generate a representation of the oculomotor space by multiplicatively combining population codes of retinal (r) and oculomotor (c) signals. LIP activity is also driven by recurrent

lateral connectivity and by feedback signals from the FEF layer. The activity of each LIP neuron is updated according to the following equation:

$$\frac{dx_i^{\text{LIP}}}{dt} = f \left(G(r)_i + \sum_j w_{ij} x_j^{\text{LIP}} + \sum_l q_{il} x_l^{\text{FEF}} - \zeta x_i^{\text{LIP}} \right) G(c)_i \quad (1)$$

where f is the logistic function, $G(r)_i$ is the retinal signal, x_j^{LIP} is the activity of other LIP neurons received through the w_{ij} lateral connection weights, x_l^{FEF} is the activity of FEF neurons that is received through the q_{il} connection weights, and ζ is the neuron's passive decay of activation (set to 0.1). $G(c)_i$ is the oculomotor signal, which represents the CD of the current eye movement, and its multiplicative interaction with the other signals produces the typical gain modulation of a BF map. Visual and motor tuning curves are Gaussian functions of r and c , respectively, both defined in retinal coordinates:

$$G(r)_i = e^{-\frac{(r-r_i)^2}{2\sigma^2}} \quad (2)$$

$$G(c)_i = e^{-\frac{(c-c_i)^2}{2\sigma^2}} \quad (3)$$

where r_i and c_i , which are uniformly spread between -40° and $+40^\circ$ in increments of 4° , indicate the preferred

retinal location and preferred oculomotor command of neuron i , respectively. The width of the Gaussians, σ , is kept fixed to 5° . Intraparietal neurons, indeed, have narrow spatial tuning, with a modal response field radius smaller than 5° (Platt & Glimcher, 1998).

Recurrent connectivity within LIP implements an internal forward model of how visual information is modified as a result of saccadic movements. To achieve this computation, the neuron with preferred target location r and oculomotor command c connects neurons with preferred retinal location and oculomotor command close to $r - c$ and 0° , respectively. The strength of these connections follows a Gaussian profile centered on $(r - c, 0^\circ)$. The lateral weight w_{ij} that connects the presynaptic neuron i to the postsynaptic neuron j within the LIP map is given by

$$w_{ij} = k \left(\theta_i e^{-\frac{((r_i - c_i) - r_j)^2 + c_j^2}{2\sigma^2}} + \rho_{ij} \right) - \varphi \quad (4)$$

where κ and φ determine the strength of excitatory and inhibitory connections, respectively ($\kappa = 3.6$; $\varphi = 0.75$). θ_i modulates the connection weight as a function of the preferred motor command of the presynaptic neuron i

$$\theta_i = -\alpha e^{-\frac{(c_i^2)}{\sigma^2}} + \beta \quad (5)$$

where α and β are set to 2.5 and 1.9, respectively. This allows sustaining memory activity during remapping and gradually decreasing it as remapping approaches its completion without altering the computational properties of the internal model. Finally, ρ_{ij} ensures that all neurons with the same preferred target location have reciprocal excitatory connections

$$\rho_{ij} = e^{-\left(-\text{abs}(c_i) - \frac{(r_i - r_j)^2}{2\sigma^2} \right)} \quad (6)$$

LIP neurons are bidirectionally connected with FEF neurons. The LIP neuron with preferred target location r is connected to FEF neurons with preferred motor command m close to r . The strength of the connection q_{il} between the LIP neuron i and the FEF neuron l is given by a Gaussian function of the difference between r_i and m_l

$$q_{il} = e^{-\frac{(r_i - m_l)^2}{2\sigma^2}} - \varphi \quad (7)$$

where φ (set to 0.75) determines the value of the inhibitory component of the connections.

The FEF map is composed of 21 neurons that generate a motor map for saccadic planning. Similar to the LIP layer, motor fields' centers of FEF neurons are uniformly

spread between -40° and $+40^\circ$ in increments of 4° . The activity of each FEF neuron is updated according to the following equation:

$$\frac{dx_i^{\text{FEF}}}{dt} = f \left(G(m)_i + \sum_n v_{in} x_n^{\text{FEF}} + \sum_l q_{il} x_l^{\text{LIP}} - 2\zeta x_i^{\text{FEF}} \right) \quad (8)$$

where f is the logistic function, x_n^{FEF} is the activity of other FEF neurons received through the v_{in} lateral connection weights, x_l^{LIP} is the activity of LIP neurons that is received through the q_{il} connection weights, and ζ is the neuron's passive decay of activation (set to 0.1). $G(m)_i$ is a Gaussian function of the planned oculomotor program, which is always set to 0 except in the case of saccadic planning and attention orienting.

Lateral connections in FEF generate local excitation and long-range inhibition. The strength of the lateral connection v_{ij} from the FEF neurons i to the FEF neuron j depends on the difference between their motor fields' centers, m_i and m_j , according to the following equation:

$$v_{ij} = \tau e^{-\frac{(m_i - m_j)^2}{2\sigma^2}} - \lambda e^{-\frac{(m_i - m_j)^2}{2\varsigma\sigma^2}} \quad (9)$$

where τ and λ determine the values of excitatory and inhibitory region and ς controls the width of the inhibitory region ($\tau = 1.15$, $\lambda = 0.47$, $\varsigma = 10$).

Continuous time was discretized in the simulations and the time constant (dt) was set to 0.01 for all simulations.

RESULTS

Ocular Perturbation Task

Before simulating attention tasks, we tested the ability of the model in performing spatial remapping by implementing a saccadic task that required to foveate a remembered spatial location after an ocular perturbation (usually evoked by electrical stimulation of the SC). Each trial started with presentation of a random visual target (r). After its offset, we simulated an ocular perturbation by generating a random CD signal (c). We decoded FEF activity using the center of mass method (Zemel, Dayan, & Pouget, 1998) and measured the error of the system (i.e., distance between expected and decoded target location) when the difference in decoded target location between two successive states was less than 0.005° (i.e., when the network has settled into a stable state). We computed the root mean square error as performance index over 300 runs with random values of r and c . The root mean square error (1.9°) was less than half of the interpeak distance in FEF, indicating that the model accurately planned the movement required to acquire the remembered target location after ocular perturbation.

The analysis of the response properties of simulated LIP neurons showed that retinotopic representations

were remapped in the coordinates of the new fixation point after ocular perturbation. Figure 2 shows the temporal evolution of the network activity. When a visual target is briefly presented to the model, a two-dimensional pattern of activity builds up in the LIP map. The hill of activity is centered at the corresponding position along the retinal axis and at 0° (corresponding to fixation) along the motor axis. After ocular perturbation, the CD signal modulates the activity in LIP, recruiting those neurons that are selective for the corresponding motor vector. Then the lateral connections, which implement the internal forward model, start to transfer the activity to the fixation neurons that code the remapped location along the retinal axis. As a result, the remapped representations in LIP are coded by those neurons whose visual RFs would have been stimulated if the visual target had still been present. This is consistent with the finding that many LIP neurons code for impending saccades (e.g., Colby, Duhamel, & Goldberg, 1996). FEF activity at the end of the remapping process encodes target position in the coordinates of the new fixation point.

Spatial Cueing Paradigm

The premotor theory maintains that motor planning generates top-down signals that produce a processing facilitation for stimuli located in the region of space toward which the motor plan was prepared. We tested this basic claim by implementing a spatial cueing paradigm (Posner,

1980), which requires detecting a visual target as fast as possible. In endogenous cueing, participants voluntarily orient to the spatial location indicated by a cognitive cue, and the target can be presented at the cued location (valid trials) or at a different location (invalid trials). In neutral condition, the cognitive cue does not indicate where to orient attention. Typically, valid trials give rise to faster RTs with respect to neutral trials (attentional benefits), whereas invalid trials give rise to slower RTs (attentional costs).

We simulated attention orienting by generating a saccadic plan in the FEF map and feeding back the activity to the LIP map. The saccadic plan could be directed toward one of two spatial locations (-4° and 4° eccentricity), similar to the classical spatial cueing paradigm (Posner, 1980). After a random delay (within the range of 300–600 cycles), we presented a visual target in the location corresponding to the planned saccade (valid condition) or in the other location (invalid condition). To measure attentional benefits and costs, we included a baseline condition in which attention orienting did not precede target presentation. We measured the number of cycles required for reaching the threshold value of 0.7 in FEF (the same response criterion was used in all subsequent simulations) as an index of RT for target detection. The target remained on until the end of the trial because, with the current set of parameters, this allowed proper build-up of activation in FEF to reach response threshold. We performed 10 runs with 60 trials each (20 valid trials, 20 invalid trials, and 20 neutral trials).





Figure 2. Snapshots of the temporal dynamics of simulated LIP (bottom) and FEF (top) neurons during the ocular perturbation task. (A) Activity profiles immediately after the onset of a visual target presented at 2° on the left from fixation. (B, C) After an ocular perturbation, which shifts the eye 18° to the left, the corresponding CD signal is combined with the memory trace of the target location. (D, E) The lateral connectivity in LIP, which implements the internal forward model, drives the hill of activity to the remapped location (16° to the right on the retinal axis and fixation position on motor axis). Note that activity in FEF reflects the changes in LIP activity. The bottom panel represents the detailed temporal sequence of events during the trial, with letters on the lower line indicating the exact timing of the activity snapshots.

Figure 3. Distance effect. Attentional costs increase as a function of distance from the cued position. (A) Simulation data. (B) Experimental data (Umiltà et al., 1994). (C) Schematic representation of the possible target positions in the simulated spatial cueing paradigm.



A repeated measures ANOVA on mean RTs showed a significant main effect of condition (valid, invalid, baseline) [$F(2, 27) = 15474, p < .0001$]. The valid condition produced faster responses than the baseline condition [265 vs. 378 cycles; $t(9) = 216.12, p < .0001$], which in turn produced faster responses than the invalid condition [378 vs. 399 cycles; $t(9) = 24.27, p < .0001$], indicating robust attentional effects for selected spatial locations in the absence of eye movements. The RT benefit observed for valid trials depends on the spatial correspondence between top-down signals (from FEF) and bottom-up signals (from the visual target) in the LIP map. In contrast, top-down and bottom-up signals are spatially misaligned during invalid trials, thereby generating two different hills of activity in LIP. The competition between these population codes through lateral connectivity slows down target detection and is responsible for the incurred RT cost.

Behavioral studies have shown that attentional costs increase as a function of the distance between target and cued location (Umiltà, Mucignat, Riggio, Barbieri, & Rizzolatti, 1994). This distance effect was attributed to the time required to reorient attention from the cued location after target presentation. To investigate the presence of a distance effect in the model, we repeated the previous simulations by adding two peripheral positions (-8° and 8°). This allowed presenting the target at four

different distances from the cued location ($4^\circ, 8^\circ, 12^\circ$, and 16°), as in the study of Umiltà and colleagues (see Figure 3B and C). A repeated measures ANOVA on mean RTs with distance ($0^\circ, 4^\circ, 8^\circ, 12^\circ$, and 16°) as factor yielded a significant main effect [$F(5, 54) = 14243, p < .0001$]. Planned comparisons (two-tailed t tests) revealed that the attentional cost varied reliably as a function of the distance from the cued location (all $ps < .0001$; see Figure 3A). Notably, the distance effect in the model emerges from lateral connectivity that generates local excitation and long-range inhibition, without requiring any additional mechanism.

Gaze-contingent Paradigm

To investigate the role of spatial remapping in attention orienting, we implemented a gaze-contingent paradigm similar to that used by Golomb et al. (2008). Following the initial phase of attention orienting (up to and including a 100-cycle fixed delay), we simulated an intervening saccade by generating a second saccadic plan in the FEF map. Because the intervening saccade was an overt eye movement, the corresponding CD signal was delivered to the network. After a variable delay (50, 100, 200, 300, 400, 500, or 600 cycles), we presented the detection target, which lasted until the end of the trial. The target could appear at the spatiotopic coordinates of the attended





Figure 4. Schematic example of target presentation in the simulated gaze-contingent paradigm. Given initial orienting to a cue location r and intervening eye movement c , we measured the latency for target detection at four different locations: r (retinotopic condition); $r - c$ (spatiotopic condition); $r + c$ (retinotopic control); $r - 2c$ (spatiotopic control), wherein c is the eye movement amplitude. The bottom right panel shows the same four locations as they appear on the display (i.e., spatiotopic coordinates).

Figure 5. Attentional facilitation in a gaze-contingent paradigm. Attentional facilitation is shown as the difference in RT for targets appearing in the spatiotopic and retinotopic locations compared with the control location baseline. Positive values reflect faster RTs than at control locations; error bars indicate 95% confidence intervals, and asterisks indicate values significantly different from 0 after Bonferroni correction for multiple comparisons. (A) Simulation data. Facilitation is plotted as a function of delay between saccadic planning and target onset. (B) Experimental data from Golomb et al. (2008; Saccade Task) on the left and from Golomb et al. (2011) on the right. Facilitation is plotted as a function of delay between the end of the intervening saccade and target onset.



location (spatiotopic condition), at its retinotopic coordinates (retinotopic condition), or in two control locations (see Figure 4) which were chosen to be equidistant from the cued position both in retinotopic and in spatiotopic coordinates. Note that the greater eccentricity of control positions has no effect in the simulations because modulation of visual acuity by eccentricity is not implemented in the model. We measured the number of cycles required for reaching the threshold value in FEF as an index of RT for target detection. To assess attentional facilitation, we computed the differences in RT when the target occurred

in the spatiotopic or retinotopic locations compared with average RT between the two control locations.

We performed 10 runs with 20 offset trials for each delay and condition. Mean RT differences were computed for each run and then entered into a repeated measures ANOVA with delay (50, 100, 200, 300, 400, 500, and 600 cycles) and condition (retinotopic vs. spatiotopic) as factors. There was no effect of condition [$F(1, 9) = 2.09, p = .18$], but there was a significant effect of delay [$F(1, 9) = 64.356, p < .0001$], and a significant interaction [$F(1, 9) = 17.82, p < .001$]. We then conducted planned




Figure 6. Snapshots of the temporal dynamics of simulated LIP (bottom) and FEF (top) neurons in the gaze contingent task. (A) Activity profiles during the initial phase of attention orienting (attention is cued to a position 8° to the right of fixation). (B) After a fixed delay of 100 cycles and the build-up of a saccadic plan in FEF (20° to the right), the corresponding CD signal is combined with the memory trace of the cued position. (C) The lateral connectivity in LIP drives the hill of activity to the remapped location. (D) After a variable delay (here 300 cycles), the target is presented in the spatiotopic (top row) or in the retinotopic position (bottom row). (E) The build-up of activity for target detection (which is achieved when activity in FEF reaches a threshold value) is facilitated in the spatiotopic trial, because the activity profile in LIP is aligned with bottom-up signals from the visual target. In contrast, top-down and bottom-up signals are spatially misaligned in the retinotopic trial, thereby generating competition between the two population codes and slowing down the RT for target detection. The bottom panel represents the detailed temporal sequence of events during the trials, with letters on the lower line indicating the exact timing of the activity snapshots. The right edge of the black bar representing the target does not indicate an offset, but the time required to reach the detection threshold, which occurred earlier in the spatiotopic trial.

t tests to compare retinotopic and spatiotopic conditions at different delays and to assess whether spatiotopic or retinotopic locations were significantly facilitated compared with the control baseline (0-cycle RT difference). All *t* tests were Bonferroni corrected for multiple comparisons ($p < .0024$) and two-tailed. The interaction depended on a different time course of facilitation between retinotopic and spatiotopic conditions (see Figure 5A). Retinotopic facilitation was strongest at the 50-cycle delay and then rapidly decreased, whereas spatiotopic facilitation reached its peak at later delays (200–300 cycles). Target detection was significantly faster at the retinotopic coordinates of the attended location until the 100-cycle delay. At this delay, retinotopic facilitation matched spatiotopic facilitation, which prevailed at longer delays (200–400 cycles).

These results are consistent with the empirical data (see Figure 5B) reported by Golomb et al. (2008, 2011), showing early facilitation effects at the eye-centered coordinates of the attended location and later benefits at its spatial coordinates. Our simulations well predict the interplay between retinotopic and spatiotopic facilitation during the first 200 msec after the eye movement (note that the number of cycles is not intended to directly map onto a millisecond scale). Figure 6 shows the temporal evolution of the network activity throughout a trial of the gaze contingent paradigm. After an eye movement, the hill of activity in LIP generated by attention orienting is shifted to the remapped location. If the delay between eye movement and target onset is sufficiently long to allow completion of the spatial updating, the activity profile in LIP becomes aligned with the bottom-up visual signal of a target presented at the spatiotopic location. Conversely, a target presented at the retinotopic location is spatially misaligned with the LIP memory activity, thereby generating a competition between the two population codes.

DISCUSSION

We examined whether a recurrent model of saccadic planning can account for attentional effects without requiring additional or specific mechanisms separate from the circuits that perform sensorimotor transformations for eye movements. The model employs BFs to simulate posterior parietal neurons involved in the representation of oculomotor space and incorporates a circuit responsible for updating remembered spatial locations across eye movements. Spatial remapping is achieved by means of horizontal connections among intraparietal neurons that implement an internal forward model of how an eye movement modifies visual information. This forward model combines the sensory inflow with the motor outflow to estimate the consequences of motor commands on the internal representation of salient locations.

Previous computational studies showed that spatial remapping can be implemented in a recurrent sensory map

by integrating an eye velocity signal (Droulez & Berthoz, 1991) or an eye position signal (Krommenhoek, Van Opstal, Gielen, & Van Gisbergen, 1993). Recurrent connections among simulated LIP neurons were used to model spatial updating in parietal cortex. However, specific connectivity and computations at the dendritic level (Quaia et al., 1998) or a dedicated memory buffer, which stores the location of one target at a time (Xing & Andersen, 2000a), were required. In contrast, sensorimotor transformations, STM, and spatial updating are handled in our model by the same computational units, which resemble the properties of posterior parietal neurons (for further discussion on the biological plausibility of the BF approach, see Pouget & Snyder, 2000). More recently, Keith and Crawford (2008) proposed a network model with feed-forward architecture that performs spatial updating by means of a lateral displacement in the hidden units' RFs. However, as noted in the Introduction, the hypothesis of shifting RFs is inconsistent with the empirical data on cross-modal anticipatory responses and on the updating of remembered spatial locations in LIP (see Cavanagh et al., 2010, for a thorough discussion).

Simulations of the spatial cueing paradigm showed the typical pattern of results reported in behavioral studies with regular attentional benefits and costs. Contrary to previous computational accounts of spatial attention (Cohen, Romero, Servan-Schreiber, & Farah, 1994; Mozer, 1991; Phaf, Van der Heijden, & Hudson, 1990), the model does not require any separate subsystem (e.g., specific nodes or unspecified "bias") to generate top-down attentional effects. Indeed, attentional facilitation depends only on feedback effects from premotor neurons to parietal neurons located downstream. Of course this demonstration does not rule out the possibility that other types of attentional mechanisms may also exist in the brain.

In addition to simulate attentional orienting in absence of eye movements, we implemented a gaze-contingent paradigm in which an eye shift is interposed between attentional allocation and target presentation. The model predicts that, after an eye movement, visuospatial attention is remapped in the coordinate of the new fixation point without requiring top-down reorienting signals. This automatic updating takes time and the native attentional code in retinotopic coordinates persists around the time of the eye movement. Indeed, simulations showed a processing facilitation at the retinotopic coordinates of the attended location immediately after an intervening saccade. As retinotopic facilitation decreases, spatiotopic facilitation increases and prevails at longer delays. These results are consistent with recent empirical studies devoted to investigating the allocation of spatial attention across eye movements (Golomb et al., 2008, 2011; Golomb, Nguyen-Phuc, et al., 2010; Golomb, Pulido, et al., 2010; Mathôt & Theeuwes, 2010).

It has to be noted that Golomb and colleagues (2008) failed to observe spatiotopic facilitation when participants were asked to retain a location in retinotopic coordinates.




Figure 7. Revision of the premotor theory of spatial attention. The premotor theory maintains that preparation of an eye movement produces a processing facilitation on neurons located upstream by means of topographic projections from FEF neurons. This allows the brain to select salient spatial locations (A). If an eye movement intervenes (B), an oculomotor signal coming from SC neurons modulates the activity in posterior parietal maps. As a result, selected locations are remapped in the coordinates of the new fixation point (C). Red patches indicate the position of fixation and black patches indicate the selected spatial location within each retinotopically organized map.

Building on this result, they argued that the updating of spatial attention occurs only when its spatiotopic coordinates are task relevant. However, other recent studies challenge this conclusion (Howe, Drew, Pinto, & Horowitz, 2011; Rolfs, Jonikaitis, Deubel, & Cavanagh, 2011). In particular, Howe and colleagues demonstrated that the attentional system automatically tracks visual objects in spatiotopic coordinates and compensates for ongoing eye movements. Moreover, Rolfs and colleagues have shown that the topography of attention is modified before a saccade to compensate for an intervening eye movement, preserving the alignment of the attentional focus with the corresponding target. Taken together, these results suggest that spatiotopic updating is automatic, although spatiotopic facilitation may be affected by task demands. More generally, the automaticity of a neural process does not necessarily imply the presence of a behavioral effect.

Our computational model represents a fundamental improvement of the premotor theory of attention, because it takes into account the mechanism responsible for updating attended locations across saccades. During execution of a saccadic movement, a CD signal of the motor command is combined with the internal representation of the attended location, which is remapped in the coordinates of the new fixation point. This allows the brain to align spatial attention with the external space, thus producing spatiotopic facilitation effects. As a result, our simulations suggest that the ability to keep attention at a spatial location while moving the eyes elsewhere is a consequence of the computations performed by parietal neurons to achieve spatial remapping. That is, the interactions between top-down orienting and spatial remapping account for behavioral dissociations between attention and eye movements that one may invoke to challenge the premotor theory. The model predicts that, while top-down selection depends on topographic projections from premotor neurons, the updating of selected locations involves an internal forward model that combine oculomotor information with visual memory signals.

The premotor theory of attention has also been questioned on the basis of a neurophysiological dissociation between attentional selection and saccadic preparation in FEF (Thompson, Biscoe, & Sato, 2005; Juan, Shorter-Jacobi, & Schall, 2004; Sato & Schall, 2003), which hinges upon the existence of two subpopulations of neurons with distinct visual and motor properties (for a review, see Awh, Armstrong, & Moore, 2006). However, it should be noted that, although visual activity in FEF does not drive saccadic-related activity, the selection of potential saccade targets by FEF visual neurons remains an essential part of saccade planning (for further discussion, see Thompson et al., 2005). Moreover, all those studies that showed a dissociation between orienting of spatial attention and saccadic preparation in FEF employed a singleton search task (see also Awh et al., 2006). This type of task is known to evoke stimulus-driven (i.e., exogenous) rather than endogenous orienting of attention. However, the premotor theory was introduced to explain endogenous orienting, and from the beginning it was made clear that it did not apply to exogenous orienting (e.g., Rizzolatti et al., 1994). Thus, these results do not invalidate the premotor theory but reinforce a fundamental distinction between endogenous and exogenous orienting, which is also endorsed by the broader model of attention orienting proposed by Corbetta and Shulman (2002).

From our revision of the premotor theory (see Figure 7), it emerges that spatial attention does not merely reflect the consequences of oculomotor preparation (covert orienting), but also the outcome of an internal dynamic estimate of how a saliency map of the visual world is modified as a result of oculomotor action (attention remapping).

In conclusion, the model provides new insights into how spatial remapping may be implemented in parietal cortex and offers a computational framework for recent proposals that link visual stability with remapping of attention pointers (Cavanagh et al., 2010). The updating of attended locations in parietal spatial maps may contribute to

the perception of a stable visual world despite continuous changes in retinal representations across eye movements.

Acknowledgments

This study was supported by grant 210922 from the European Research Council to M. Z.

Reprint requests should be sent to Marco Zorzi, Dipartimento di Psicologia Generale, Università di Padova, via Venezia 8, 35131 Padova, Italy, or via e-mail: marco.zorzi@unipd.it.

REFERENCES

- Andersen, R. A. (1989). Visual and eye movement functions of the posterior parietal cortex. *Annual Review of Neuroscience*, *12*, 377–403.
- Andersen, R. A., Essick, G. K., & Siegel, R. M. (1985). Encoding of spatial location by posterior parietal neurons. *Science*, *230*, 456–458.
- Astle, D. E. (2009). Going from a retinotopic to a spatiotopic coordinate system for spatial attention. *Journal of Neuroscience*, *29*, 3971–3973.
- Awh, E., Armstrong, K., & Moore, T. (2006). Visual and oculomotor selection: Links, causes and implications for spatial attention. *Trends in Cognitive Sciences*, *10*, 124–130.
- Awh, E., & Jonides, J. (2001). Overlapping mechanisms of attention and spatial working memory. *Trends in Cognitive Sciences*, *5*, 119–126.
- Beauchamp, M. S., Petit, L., Ellmore, T. M., Ingeholm, J., & Haxby, J. V. (2001). A parametric fMRI study of overt and covert shifts of visuospatial attention. *Neuroimage*, *14*, 310–321.
- Bisley, J. W., & Goldberg, M. E. (2003). Neuronal activity in the lateral intraparietal area and spatial attention. *Science*, *299*, 81–86.
- Brozovic, M., Abbott, L. F., & Andersen, R. A. (2008). Mechanisms of gain modulation at single neuron and network levels. *Journal of Computational Neuroscience*, *25*, 158–168.
- Cavanagh, P., Hunt, A. R., Afraz, A., & Rolfs, M. (2010). Visual stability based on remapping of attention pointers. *Trends in Cognitive Sciences*, *14*, 147–153.
- Chambers, C. D., & Mattingley, J. B. (2005). Neurodisruption of selective attention: Insights and implications. *Trends in Cognitive Sciences*, *11*, 542–550.
- Clower, D. M., West, R. A., Lynch, J. C., & Strick, P. L. (2001). The inferior parietal lobule is the target of output from the superior colliculus, hippocampus, and cerebellum. *Journal of Neuroscience*, *21*, 6283–6291.
- Cohen, J. D., Romero, R. D., Servan-Schreiber, D., & Farah, M. J. (1994). Mechanisms of spatial attention: The relation of macrostructure to microstructure in parietal neglect. *Journal of Cognitive Neuroscience*, *6*, 377–387.
- Colby, C. L., Duhamel, J. R., & Goldberg, M. E. (1996). Visual, presaccadic, and cognitive activation of single neurons in monkey lateral intraparietal area. *Journal of Neurophysiology*, *76*, 2841–2852.
- Colby, C. L., & Goldberg, M. E. (1999). Space and attention in parietal cortex. *Annual Review of Neuroscience*, *22*, 319–349.
- Compte, A., Brunel, N., Goldman-Rakic, P. S., & Wang, X. J. (2000). Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cerebral Cortex*, *10*, 910–923.
- Corbetta, M., Akbudak, E., Conturo, T. E., Snyder, A. Z., Ollinger, J. M., Drury, H. A., et al. (1998). A common network of functional areas for attention and eye movements. *Neuron*, *21*, 761–773.
- Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, *3*, 201–215.
- Denève, S., Duhamel, J., & Pouget, A. (2007). Optimal sensorimotor integration in recurrent cortical networks: A neural implementation of Kalman filters. *Journal of Neuroscience*, *27*, 5744–5756.
- Deneve, S., Latham, P., & Pouget, A. (2001). Efficient computation and cue integration with noisy population codes. *Nature Neuroscience*, *4*, 826–831.
- Desmurget, M., & Grafton, S. (2000). Forward modeling allows feedback control for fast reaching movements. *Trends in Cognitive Sciences*, *4*, 423–431.
- Droulez, J., & Berthoz, A. (1991). A neural network model of sensoritopic maps with predictive short-term memory properties. *Proceedings of the National Academy of Sciences, U.S.A.*, *88*, 9653–9657.
- Duhamel, J. R., Colby, C. L., & Goldberg, M. E. (1992). The updating of the representation of visual space in parietal cortex by intended eye movements. *Science*, *255*, 90–92.
- Ferraina, S., Pare, M., & Wurtz, R. H. (2002). Comparison of cortico-cortical and cortico-collicular signals for the generation of saccadic eye movements. *Journal of Neurophysiology*, *87*, 845–858.
- Golomb, J. D., Chun, M. M., & Mazer, J. (2008). The native coordinate system of spatial attention is retinotopic. *Journal of Neuroscience*, *28*, 10654–10662.
- Golomb, J. D., Marino, A. C., Chun, M. M., & Mazer, J. A. (2011). Attention doesn't slide: Spatiotopic updating after eye movements instantiates a new, discrete attentional locus. *Attention, Perception, and Psychophysics*, *73*, 7–14.
- Golomb, J. D., Nguyen-Phuc, A. Y., Mazer, J. A., McCarthy, G., & Chun, M. M. (2010). Attentional facilitation throughout human visual cortex lingers in retinotopic coordinates after eye movements. *Journal of Neuroscience*, *30*, 10493–10506.
- Golomb, J. D., Pulido, V. Z., Albrecht, A. R., Chun, M. M., & Mazer, J. A. (2010). Robustness of the retinotopic attentional trace after eye movements. *Journal of Vision*, *10*, 19.1–19.12.
- Howe, P. D. L., Drew, T., Pinto, Y., & Horowitz, T. S. (2011). Remapping in multiple object tracking. *Vision Research*, *51*, 489–495.
- Ignashchenkova, A., Dicke, P. W., Haarmeier, T., & Thier, P. (2004). Neuron-specific contribution of the superior colliculus to overt and covert shifts of attention. *Nature Neuroscience*, *7*, 56–64.
- Juan, C.-H., Shorter-Jacobi, S. M., & Schall, J. D. (2004). Dissociation of spatial attention and saccade preparation. *Proceedings of the National Academy of Sciences, U.S.A.*, *101*, 15541–15544.
- Kawato, M. (1999). Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology*, *9*, 718–727.
- Keith, G. P., & Crawford, J. D. (2008). Saccade-related remapping of target representations between topographic maps: A neural network study. *Journal of Computational Neuroscience*, *24*, 157–178.
- Krommenhoek, K. P., Van Opstal, A. J., Gielen, C. C., & Van Gisbergen, J. A. (1993). Remapping of neural activity in the motor colliculus: A neural network study. *Vision Research*, *33*, 1287–1298.
- Kustov, A. A., & Robinson, D. L. (1996). Shared neural control of attentional shifts and eye movements. *Nature*, *384*, 74–77.
- Mathôt, S., & Theeuwes, J. (2010). Evidence for the predictive remapping of visual attention. *Experimental Brain Research*, *200*, 117–122.

- Mesulam, M. M. (1990). Large-scale neurocognitive networks and distributed processing for attention, language, and memory. *Annals of Neurology*, *28*, 597–613.
- Moore, T., Armstrong, K. M., & Fallah, M. (2003). Visuomotor origins of covert spatial attention. *Neuron*, *40*, 671–683.
- Moore, T., & Fallah, M. (2001). Control of eye movements and spatial attention. *Proceedings of the National Academy of Sciences, U.S.A.*, *98*, 1273–1276.
- Moore, T., & Fallah, M. (2004). Microstimulation of the frontal eye field and its effects on covert spatial attention. *Journal of Neurophysiology*, *91*, 152–162.
- Mozer, M. C. (1991). *The perception of multiple objects: A connectionist approach*. Cambridge, MA: MIT Press/Bradford Books.
- Muller, J. R., Philastides, M. G., & Newsome, W. T. (2005). Microstimulation of the superior colliculus focuses attention without moving the eyes. *Proceedings of the National Academy of Sciences, U.S.A.*, *102*, 524–529.
- Nobre, A. C., Gitelman, D. R., Dias, E. C., & Mesulam, M. M. (2000). Covert visual spatial orienting and saccades: Overlapping neural systems. *Neuroimage*, *3*, 210–216.
- O'Reilly, R. C. (1998). Six principles for biologically-based computational models of cortical cognition. *Trends in Cognitive Sciences*, *2*, 455–462.
- Patel, G. H., Shulman, G. L., Baker, J. T., Akbudak, E., Snyder, A. Z., Snyder, L. H., et al. (2010). Topographic organization of macaque area LIP. *Proceedings of the National Academy of Sciences, U.S.A.*, *107*, 4728–4733.
- Perry, C., Ziegler, J. C., & Zorzi, M. (2007). Nested incremental modeling in the development of computational theories: The CDP+ model of reading aloud. *Psychological Review*, *114*, 273–315.
- Phaf, R. H., Van der Heijden, A. H., & Hudson, P. T. (1990). SLAM: A connectionist model for attention in visual selection tasks. *Cognitive Psychology*, *22*, 273–341.
- Platt, M. L., & Glimcher, P. W. (1998). Response fields of intraparietal neurons quantified with multiple saccadic targets. *Experimental Brain Research*, *121*, 65–75.
- Poggio, T. (1990). A theory of how the brain might work. *Cold Spring Harbor Symposium on Quantitative Biology*, *55*, 899–910.
- Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, *32*, 3–25.
- Pouget, A., & Sejnowski, J. T. (1997). Spatial transformations in the parietal cortex using basis functions. *Journal of Cognitive Neuroscience*, *9*, 222–237.
- Pouget, A., & Snyder, L. H. (2000). Computational approaches to sensorimotor transformations. *Nature Neuroscience*, *3*, 1192–1198.
- Quaia, C., Optician, L. M., & Goldberg, M. E. (1998). The maintenance of spatial accuracy by the perisaccadic remapping of visual receptive fields. *Neural Networks*, *11*, 1229–1240.
- Rizzolatti, G., Riggio, L., Dascola, I., & Umiltà, C. (1987). Reorienting attention across the horizontal and vertical meridians: Evidence in favor of a premotor theory of attention. *Neuropsychologia*, *25*, 31–40.
- Rizzolatti, G., Riggio, L., & Sheliga, B. M. (1994). Space and selective attention. In C. Umiltà & M. Moscovitch (Eds.), *Attention and performance XV* (pp. 231–265). Cambridge, MA: MIT Press.
- Rolfs, M., Jonikaitis, D., Deubel, H., & Cavanagh, P. (2011). Predictive remapping of attention across eye movements. *Nature Neuroscience*, *14*, 252–256.
- Salinas, E., & Abbott, L. F. (1995). Transfer of coded information from sensory to motor networks. *Journal of Neuroscience*, *15*, 6461–6474.
- Salinas, E., & Abbott, L. F. (1996). A model of multiplicative neural responses in parietal cortex. *Proceedings of the National Academy of Sciences, U.S.A.*, *93*, 11956–11961.
- Salinas, E., & Thier, P. (2000). Gain modulation: A major computational principle of the central nervous system. *Neuron*, *27*, 15–21.
- Sato, T. R., & Schall, J. D. (2003). Effects of stimulus-response compatibility on neural selection in frontal eye field. *Neuron*, *38*, 637–648.
- Smith, M. A., & Crawford, J. D. (2005). Distributed population mechanism for the 3-D oculomotor reference frame transformation. *Journal of Neurophysiology*, *93*, 1742–1761.
- Sommer, M. A., & Wurtz, R. H. (2008). Brain circuits for the internal monitoring of movements. *Annual Review of Neuroscience*, *31*, 317–338.
- Thompson, K. G., Biscoe, K. L., & Sato, T. R. (2005). Neuronal basis of covert spatial attention in the frontal eye field. *Journal of Neuroscience*, *25*, 9479–9487.
- Todorov, E. (2004). Optimality principles in sensorimotor control. *Nature Neuroscience*, *7*, 907–915.
- Umiltà, C., Mucignat, C., Riggio, L., Barbieri, C., & Rizzolatti, G. (1994). Programming shifts of spatial attention. *European Journal of Cognitive Psychology*, *6*, 23–41.
- Umiltà, C., Riggio, L., Dascola, I., & Rizzolatti, G. (1991). Differential effects of central and peripheral cues on the reorienting of spatial attention. *European Journal of Cognitive Psychology*, *3*, 247–267.
- Vardak, C., Ibos, G., Duhamel, J. R., & Olivier, E. (2006). Contribution of the monkey frontal eye field to covert visual attention. *Journal of Neuroscience*, *26*, 4228–4235.
- Vaziri, S., Diedrichsen, J., & Shadmehr, R. (2006). Why does the brain predict sensory consequences of oculomotor commands? Optimal integration of the predicted and actual sensory feedback. *Journal of Neuroscience*, *26*, 4188–4197.
- Wang, X. J. (2001). Synaptic reverberation underlying mnemonic persistent activity. *Trends in Neurosciences*, *24*, 455–463.
- Wolpert, D. M., Ghahramani, Z., & Jordan, M. (1995). An internal model for sensorimotor integration. *Science*, *269*, 1880–1882.
- Xing, J., & Andersen, R. A. (2000a). Memory activity of LIP neurons for sequential eye movements simulated with neural networks. *Journal of Neurophysiology*, *84*, 651–665.
- Xing, J., & Andersen, R. A. (2000b). Models of the posterior parietal cortex which perform multimodal integration and represent space in several coordinate frames. *Journal of Cognitive Neuroscience*, *12*, 601–614.
- Zemel, R. S., Dayan, P., & Pouget, A. (1998). Probabilistic interpretation of population codes. *Neural Computing*, *10*, 403–430.
- Zipser, D., & Andersen, R. A. (1988). A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature*, *331*, 679–684.