Behavioral/Systems/Cognitive

An Integrated Microcircuit Model of Attentional Processing in the Neocortex

Salva Ardid,[1,2] **Xiao-Jing Wang,**[3] **and Albert Compte**[1,2]

[1]Instituto de Neurociencias de Alicante, Universidad Miguel Hernández–Consejo Superior de Investigaciones Científicas, 03550 Sant Joan d'Alacant, Spain, [2]Institut d'Investigacions Biomèdiques August Pi i Sunyer, 08036 Barcelona, Spain, and [3]Department of Neurobiology and Kavli Institute for Neuroscience, Yale University School of Medicine, New Haven, Connecticut 06520-8001

Selective attention is a fundamental cognitive function that uses top-down signals to orient and prioritize information processing in the brain. Single-cell recordings from behaving monkeys have revealed a number of attention-induced effects on sensory neurons, and have given rise to contrasting viewpoints about the neural underpinning of attentive processing. Moreover, there is evidence that attentional signals originate from the prefrontoparietal working memory network, but precisely how a source area of attention interacts with a sensory system remains unclear. To address these questions, we investigated a biophysically based network model of spiking neurons composed of a reciprocally connected loop of two (sensory and working memory) networks. We found that a wide variety of physiological phenomena induced by selective attention arise naturally in such a system. In particular, our work demonstrates a neural circuit that instantiates the "feature-similarity gain modulation principle," according to which the attentional gain effect on sensory neuronal responses is a graded function of the difference between the attended feature and the preferred feature of the neuron, independent of the stimulus. Furthermore, our model identifies key circuit mechanisms that underlie feature-similarity gain modulation, multiplicative scaling of tuning curve, and biased competition, and provide specific testable predictions. These results offer a synthetic account of the diverse attentional effects, suggesting a canonical neural circuit for feature-based attentional processing in the cortex.

Key words: feature-based attention; cortical circuits; working memory; sensory systems; computational model; top-down; control

Introduction

In cluttered environments, efficient vision depends critically on appropriate allocation and deployment of voluntary attention (Desimone and Duncan, 1995; Colby and Goldberg, 1999; Itti and Koch, 2001). Neurobiologists have identified different forms of attentional modulation of neuronal responses in the sensory cortex, notably in visual areas V4 and visual middle-temporal area (MT), and this has led to the formulation of different conceptual models for the physiological action of attention.

Biased competition (BC) (Desimone and Duncan, 1995) holds that attention biases the competition between the representations of stimuli within the receptive field (RF) of the neuron. This is based on experiments showing that neural responses to simultaneous presentation of preferred and nonpreferred stimuli in the RF are intermediate between responses elicited by either stimulus alone, and attention focused onto one of the stimuli leads to responses determined by the attended stimulus, as if the

unattended stimulus was absent (Moran and Desimone, 1985; Desimone and Duncan, 1995; Treue and Maunsell, 1996; Reynolds et al., 1999).

Other researchers have proposed the multiplicative gain modulation interpretation (MGM). This is based on experiments showing how the gain of neuronal responses to a single stimulus is enhanced both when attention is focused inside the RF of the neuron (McAdams and Maunsell, 1999) or onto the preferred feature of the neuron (Treue and Martinez-Trujillo, 1999), without affecting the selectivity of the cell.

And, recently, it has been observed in MT that attention modulates population activity by enhancing it at the attentional focus and suppressing it in its surround (Martinez-Trujillo and Treue, 2004). This finding can be understood within the feature-similarity gain principle (FSGP) (Treue and Martinez-Trujillo, 1999; Martinez-Trujillo and Treue, 2004; Boynton, 2005; Maunsell and Treue, 2006), which posits that selective attention to a given feature modulates the firing rate of a neuron by a gain factor, which depends on the parametrical similarity between the attended feature (spatial location, orientation, direction. . .) and the preference of the neuron for that feature. Neurons multiply their firing whenever their preferences are parametrically close to the attended feature; otherwise, neural responses are divided, resulting in a population selectivity enhancement. In contrast with BC, this principle emphasizes that attention modifies neural responses multiplicatively, with the same factor regardless of sensory inputs (Maunsell and Treue, 2006). Physiologically based

computational models have been proposed for BC (Deco and Rolls, 2005); in contrast, FSGP has been considered only conceptually (Boynton, 2005). How the precise algorithmic formulation of attentional modulation in FSGP could be instantiated neurophysiologically remains unclear.

Because these conceptual models emphasize contrasting computational algorithms that apply to both spatial and featural attention, it remains unclear whether a unified mechanistic framework can account for all of these experimental results. We addressed this question using a biophysically plausible computational model of two interacting cortical networks of spiking neurons that integrates the source area of the attentional top-down signal. We modeled the visual area MT, selective to direction of motion, and we assumed that it received a top-down attentional signal originating in a working memory area (Desimone and Duncan, 1995; Barceló et al., 2000; Hopfinger et al., 2000; Miller and Cohen, 2001; Corbetta and Shulman, 2002; Lebedev et al., 2004; Hagler and Sereno, 2006; Grent-t Jong and Woldorff, 2007). Note that the hypothesized role of a working memory circuit as a source area of attentional signals still awaits explicit experimental validation in attention-to-motion tasks. In this work, we tested the idea that the interactions between these two areas are sufficient to explain the above-described attentional effects.

In addition to reproducing the experimental observations, our model is the first explicit neurophysiological implementation of FSGP. And, together, it proves the compatibility between FSGP and BC (Boynton, 2005; Deco and Rolls, 2005), assigning them, for the first time, specific (and dissectable) mechanisms within the cortical circuitry. Such a comprehensive account of neurophysiological data of attention in a biophysical network model leads us to suggest that it constitutes the backbone of a "canonical" neural circuit for feature-based attentional processing.

Materials and Methods

The model network. Each of the two network modules represents a local circuit of the cortex. The sensory network represents a local circuit of the MT, and we refer to the working memory module as a local circuit of the prefrontal cortex (PFC) for the sake of simplicity, although working memory and selective attention are likely to be subserved by both prefrontal and parietal cortices (Colby and Goldberg, 1999; Hopfinger et al., 2000; Corbetta and Shulman, 2002; Lebedev et al., 2004; Moore, 2006; Grent-t Jong and Woldorff, 2007). In addition, there is anatomical evidence of reciprocal connections between PFC and MT (Barbas, 1988; Schall et al., 1995; Burman et al., 2006). The MT and PFC circuits had exactly the same wiring structure; they only differed in the strength of the synaptic connectivity within each module: the PFC module had strong recurrent excitatory connections to sustain persistent activity, whereas MT was dominated by inhibition. A detailed account of the local circuit model can also be found in the study by Compte et al. (2000). For each circuit, pyramidal cells ($N_E = 1024$) and interneurons ($N_I = 256$) were spatially distributed on a ring simulating the cortical columnar organization, labeled by their preferred direction of motion (θ_{pref}, from 0 to 360°). Their axonal collaterals differentially targeted neighboring (isodirectional) and distant (crossdirectional) neurons. This was implemented by taking the synaptic conductance between neuron i and neuron j to be $g_{syn,ij} = W(\theta_i - \theta_j)G_{syn}$, where $W(\theta_i - \theta_j)$ was either a constant for unstructured connections ($W(\theta_i - \theta_j) = 1$), or the sum of a constant term plus a Gaussian: $W(\theta_i - \theta_j) = J^- + (J^+ - J^-) \exp(-(\theta_i - \theta_j)^2/2\sigma^2)$. In both networks, only the excitatory-to-excitatory connectivity was structured with $\sigma_{EE} = 14.4°$ and $J_{EE}^+ = 1.62$ (Compte et al., 2000). The excitatory-to-inhibitory, inhibitory-to-excitatory, and inhibitory-to-inhibitory connections were unstructured (i.e., the cross-directional and isodirectional components of feedback inhibitory connections were

equally strong). Following the notations in the study by Compte et al. (2000), the parameters defining the strengths of local connections in the two networks were as follows: in PFC, $G_{EE,AMPA} = 0.391$ nS, $G_{EE,NMDA} = 0.732$ nS (pyramid-to-pyramid); $G_{EI,AMPA} = 0.293$ nS, $G_{EI,NMDA} = 0.566$ nS (pyramid-to-interneuron); $G_{IE} = 3.74$ nS (interneuron-to-pyramid); $G_{II} = 2.87$ nS (interneuron-to-interneuron); in MT: $G_{EE,AMPA} = 0.005$ nS, $G_{EE,NMDA} = 0.093$ nS (pyramid-to-pyramid); $G_{EI,AMPA} = 0.005$ nS, $G_{EI,NMDA} = 0.195$ nS (pyramid-to-interneuron); $G_{IE} = 1.47$ nS (interneuron-to-pyramid); $G_{II} = 0.391$ nS (interneuron-to-interneuron). Thus, recurrent excitation was between 1 and 2 orders of magnitude stronger in PFC than in MT, and synaptic inhibition was very strong in both modules.

Both pyramidal cells and interneurons were modeled as leaky integrate-and-fire neurons, with the same parameters as for neurons in the network model of Compte et al. (2000). Specifically, each type of cell was characterized by six intrinsic parameters: the total capacitance, C_m; the total leak conductance, g_L; the leak reversal potential, E_L; the threshold potential, V_{th}; the reset potential, V_{res}; and the refractory time, τ_{ref}. The values used were as follows: $C_m = 0.5$ nF, $g_L = 25$ nS, $E_L = -70$ mV, $V_{th} = -50$ mV, $V_{res} = -60$ mV, and $\tau_{ref} = 2$ ms for pyramidal cells; and $C_m = 0.2$ nF, $g_L = 20$ nS, $E_L = -70$ mV, $V_{th} = -50$ mV, $V_{res} = -60$ mV, and $\tau_{ref} = 1$ ms for interneurons. All cells received random background excitatory inputs. This unspecific external input was modeled as uncorrelated Poisson spike trains to each neuron at a rate of $\nu_{ext} = 1800$ Hz per cell (or equivalently, 1000 presynaptic Poisson spike trains at 1.8 Hz). This input was exclusively mediated by AMPA receptors (AMPARs), with the maximum conductances $g_{ext,E} = 3.1$ nS on pyramidal cells and $g_{ext,I} = 2.38$ nS on interneurons in PFC; and $g_{ext,E} = 15$ nS and $g_{ext,I} = 4.5$ nS in MT.

Neurons received their recurrent excitatory inputs through AMPAR- and NMDA receptor (NMDAR)-mediated transmission and their inhibitory inputs through GABA$_A$ receptors (GABA$_A$Rs). These conductance-based synaptic responses were calibrated by the experimentally measured dynamics of synaptic currents. Thus, postsynaptic currents were modeled according to $I_{syn} = g_{syn}s(V - V_{syn})$, where g_{syn} is a synaptic conductance, s is a synaptic gating variable, and V_{syn} is the synaptic reversal potential ($V_{syn} = 0$ for excitatory synapses; $V_{syn} = -70$ mV for inhibitory synapses). AMPAR and GABA$_A$R synaptic gating variables were modeled as an instantaneous jump of magnitude 1 when a spike occurred in the presynaptic neuron followed by an exponential decay with time constant 2 ms for AMPA and 10 ms for GABA$_A$. The NMDA conductance was voltage dependent, with g_{syn} multiplied by $1/(1 + [Mg^{2+}] \exp(-0.062 V_m)/3.57)$, $[Mg^{2+}] = 1.0$ mM. The channel kinetics was modeled by the following equations:

$$\frac{ds}{dt} = -\frac{1}{\tau_s}s + \alpha_s x(1 - s) \qquad \frac{dx}{dt} = -\frac{1}{\tau_x}x + \sum_i \delta(t - t_i),$$

where s is the gating variable, x is a synaptic variable proportional to the neurotransmitter concentration in the synapse, t_i are the presynaptic spike times, $\tau_s = 100$ ms is the decay time of NMDA currents, $\tau_x = 2$ ms controls the rise time of NMDAR channels, and $\alpha_s = 0.5$ kHz controls the saturation properties of NMDAR channels at high presynaptic firing frequencies. Parameters for synaptic transmission were taken from the study by Compte et al. (2000).

The MT and PFC network modules were interconnected through bottom-up and top-down AMPAR-mediated connections (see scheme in Fig. 1). Both bottom-up and top-down connectivities were topographic, so that for both the bottom-up and top-down pathways neurons sharing the same preference were more strongly coupled than neurons with disparate preferences. This connectivity was described by a Gaussian function: $g_{syn} = G_{syn} \exp(-(\theta_i - \theta_j)^2/2\sigma^2)/\sigma\sqrt{2\pi}$. We used for the bottom-up connection onto PFC pyramids, $G_{EE}^{MT \to PFC} = 0.005$ nS and $\sigma = 36°$; for the bottom-up connection onto PFC interneurons, $G_{EI}^{MT \to PFC} = 0$; for the top-down connection onto MT pyramids, $G_{EE}^{PFC \to MT} = 0.146$ nS and $\sigma = 72°$; and for the top-down connection onto MT interneurons, $G_{EI}^{PFC \to MT} = 0.039$ nS and $\sigma = 72°$.

The simulations. The simulation protocol was chosen to resemble the behavioral protocol used in the experiment of Martinez-Trujillo and

Treue (2004). There, monkeys were trained to fixate a central spot during a brief presentation of a peripheral random-dot pattern in coherent motion, which was the stimulus to be attended. Subsequently, an additional random-dot pattern in coherent motion was added in the receptive field of the neuron (test stimulus), which was behaviorally irrelevant but could share or not the direction of motion of the attended stimulus. The monkey had to report a direction or speed change in the attended stimulus and ignore changes in the test stimulus. The experiment revealed attentional modulation on the neuronal responses to the unattended test stimulus that depends on the attended direction of motion (feature-based attention). Here, motion stimulus presentation to the network was modeled through selective transient current injection to MT cells (see below, Task-related extrinsic inputs). We included a delay period or D-period (see D in Figs. 2A, 4A), between the presentation of an attended feature (cue period or C-period) (see C in Figs. 2A, 4A) and the presentation of the test stimulus (test period or T-period) (see T in Figs. 2A, 4A). During the D-period, the visual stimulus was absent. By including this period, we were able to evaluate the effect of an attentional bias on the MT network baseline activity (see Fig. 4C).

Task-related extrinsic inputs. Cells in area MT received external inputs from primary visual area V1, which were selective to the direction of motion of the visually presented stimulus (Born and Bradley, 2005). We thus modeled motion stimuli presentation by injecting external currents to MT neurons that mimicked outputs from V1 to MT. We tried with Poisson-triggered synaptic inputs and our conclusions remained unaffected. When there was a single motion direction (θ_S), the current injected to a neuron labeled by θ_i was $I(\theta_i) = I_0 + I_1 \exp(\mu(\cos(\theta_i - \theta_S) - 1))$; for MT pyramids, we used $I_0^E = 1$ nA and $I_1^E = 0.9$ nA; for MT interneurons, we used $I_0^I = 0.2$ nA and $I_1^I = 0.18$ nA; and for both cell types, $\mu = 2.53$ (this choice of μ gives a connectivity profile very close to a Gaussian with a constant baseline, with the same width as MT-to-PFC connections). When two overlapping directions of motion were visually presented, the current impinging on MT neurons was the sum of the currents corresponding to the two single stimuli, normalized so that the maximal current was still $I_0 + I_1$ (supplemental Fig. 1A, available at www.jneurosci.org as supplemental material).

This normalization was derived from the observation that the maximal response of a direction-selective V1 neuron remains the same for either single motion or transparent motion stimuli (Snowden et al., 1991). More abstract models of V1 neurons selective to motion direction typically include a similar normalizing factor (Simoncelli and Heeger, 1998; Rust et al., 2006). PFC model neurons received motion-specific sensory inputs only through the MT-to-PFC pathway.

In all our simulation trials and during the identical C-period, all PFC neurons also received a constant current injection of 0.025 nA. This current was not selective, and thus it did not carry any direction of motion information. It was too weak to trigger by itself a persistent activity pattern in the PFC network (see Fig. 2A, left), but strong enough so that, when presented coincidentally with a visual stimulus, the PFC was able to store the directional information from MT (see Fig. 2A, right). Such a "gating input" allows our model to differentiate an attentional cue from a visual stimulus presented during the T-period in Figure 2A, left.

The integration method used was a second-order Runge–Kutta algorithm with a time step of $\Delta t = 0.02$ ms. The custom code for the simulations was written in C++.

Results

The network model architecture

We built a network model of spiking neurons composed of two interacting areas, a sensory area selective for motion direction (MT) and a working memory area that selectively stored this information. The internal structure in each of the two local networks is in accordance with the known anatomical and physiological characteristics of cortical microcircuitry. The interareal reciprocal connections followed some simple rules, based also on biological plausibility: neurons with similar preferred directions were more strongly connected (following a Gaussian function),
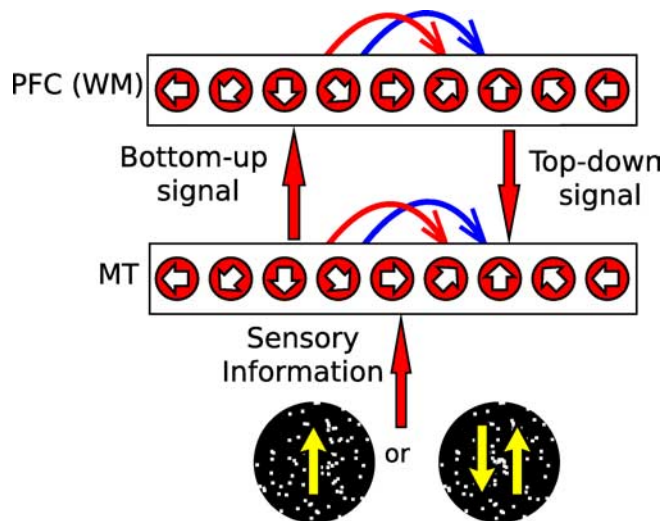

Figure 1. Scheme of the loop architecture (red is excitation, and blue is inhibition). Two kinds of motion stimuli are considered (random-dot patterns; yellow arrows indicate signal motion directions): single (left) and transparent (right) motion. WM, Working memory.

and synapses were all excitatory, but could target both pyramidal cells and interneurons. For explicit details, see Materials and Methods.

Our model was constrained based on a number of specific experimental results in area MT. On the one hand, neural responses to a motion stimulus in the receptive field have been quantitatively characterized (Maunsell and Van Essen, 1983; Snowden et al., 1992). On the other hand, there is evidence that the circuits in area MT are endowed with competition mechanisms, because the spiking response of an MT cell (but not a V1 cell) is suppressed when two superimposed moving random dot patterns are presented (Snowden et al., 1991; Treue et al., 2000) (see scheme of these stimuli in Fig. 1). We used this data to constrain our MT network model. To this end, we used both full-scale simulations (see Materials and Methods) and a mean-field approximation (Renart et al., 2003) of the MT network to allow for extensive parameter space exploration. We found that the appropriate competitive responses could be realized (supplemental Fig. 1, available at www.jneurosci.org as supplemental material) if bottom-up inputs into MT targeted also local-circuit inhibitory neurons. Interneurons project onto excitatory cells strongly to provide inhibition commensurate with the overall feedforward drive, thereby instantiating a circuit mechanism for normalization (Simoncelli and Heeger, 1998; Rust et al., 2006). The downstream working memory area was modeled as by Compte et al. (2000). This module will be referred to as prefrontal cortex (PFC) module for the sake of simplicity, although working memory and selective attention are likely to be subserved by both prefrontal and parietal cortices (Colby and Goldberg, 1999; Hopfinger and Mangun, 2000; Corbetta and Shulman, 2002; Lebedev et al., 2004; Moore, 2006; Grent-t Jong and Woldorff, 2007). In this model, working memory of a directional cue is achieved through reverberatory interactions between spiking neurons in the local network. Thus, the two networks in our model share the same qualitative internal architecture, but the PFC module is endowed with strong recurrent excitation, whereas the MT network is dominated by inhibition. Both cortical network modules are reciprocally connected with topographically specific bottom-up and top-down synaptic connections (Fig. 1) to explore the orienting effects of a selective firing pattern in the PFC
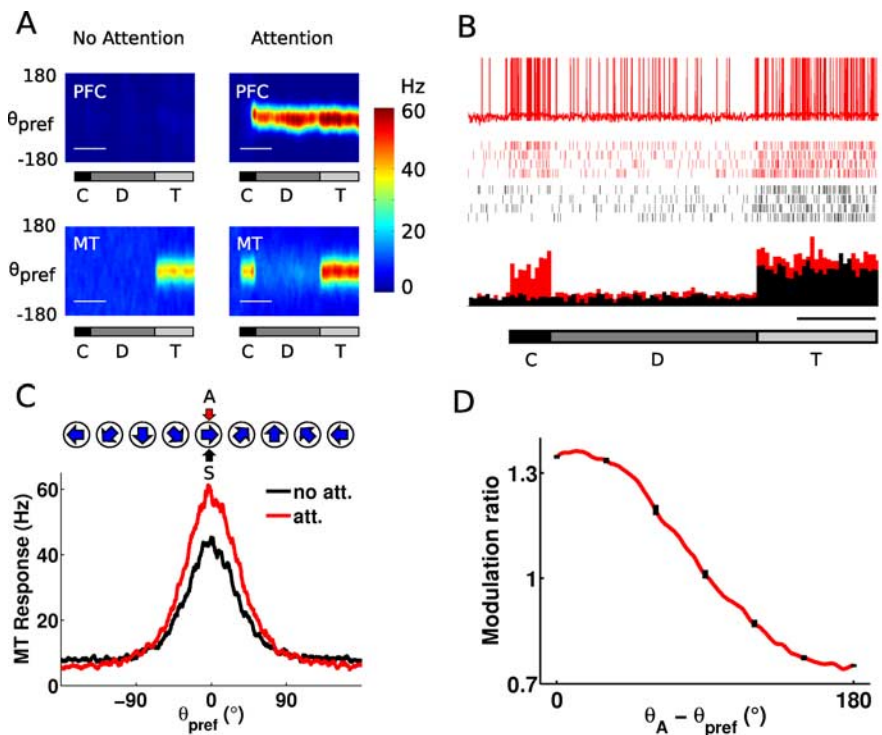
Figure 2. Feature-based attention in model simulations for single motion stimulus. *A*, Network activity for an unattended (left) and an attended trial (right). *x*-Axis, Time; *y*-axis, neurons labeled by preferred direction θ_{pref}. Activity is color-coded. C, Cue period; D, delay period; T, test period. Calibration, 1 s. *B*, Activity of a neuron with $\theta_{pref} = \theta_S$. Top, Sample membrane potential; middle, spike trains in several trials; bottom, trial-averaged activity (red, attended; black, unattended trials) (calibration: time, 1 s; voltage, 50 mV; rate, 40 Hz). *C*, Selective enhancement of MT population activity. The scheme (top) depicts how the curves (attended in red; unattended in black) were generated: for fixed test stimulus θ_S and attended feature θ_A, the activity of all neurons (blue arrows) were measured. *D*, Smoothed modulation ratio (firing rate with attention divided by that without attention) plotted against the distance between attended feature and neuronal preference.

network over the population activity and single neuron responses in the MT network. The bottom-up connection parameters were tuned to allow the transmission of visual information from the MT to the PFC module.

The parameters of the top-down connection were tuned to produce selectivity enhancement of MT neural population responses (Fig. 2C) in agreement with experimental data (Martinez-Trujillo and Treue, 2004). We will call this selective enhancement with inhibitory surround of population responses, where by inhibitory surround we mean that peak responses are enhanced and surround responses are suppressed. Such an inhibitory ring around the attentional focus has been recently validated in imaging studies in humans as well (Hopf et al., 2006). The selectivity enhancement of population profiles is a relevant finding, because what matters functionally are instantaneous population activity patterns rather than neuronal tuning curves obtained from multiple trials.

The rest of the phenomenology reported here (see Figs. 3–5) emerged from the model constrained this way, without any further parameter tuning.

Attentional enhancement of population selectivity

A single-trial simulation (Fig. 2A, right) consisted of three task epochs. In a cue period (C), a transient input about attended motion direction (θ_A) triggered a self-sustained persistent activity (peaked at the attended directional angle θ_A) in the PFC network. In simulations, it was done by a combination of a directional stimulus (θ_A) to the MT network, whose activity was

projected through weak bottom-up connections to PFC, and a transient nonspecific input to the PFC module (see Materials and Methods). This "gating input" was weak enough so it did not trigger by itself persistent activity in the PFC module (Fig. 2A, left), but it allowed this activity to develop in case a stimulus was simultaneously presented to the MT network (Fig. 2A, right). A plausible physiological substrate for this input could be found in the phasic alertness circuits recently identified in the superior temporal gyrus or in the thalamus (Sturm and Wilmes, 2001; Fan et al., 2005; Thiel and Fink, 2007). We thus assume that projections from these areas generate a slightly net increase of unspecific external input to PFC neurons during the C-period of our task. The C-period was followed by a delay period (D), where in the absence of all external inputs PFC maintained the information of the attended feature, if presented. Finally, in a test period (T), a test stimulus θ_S was presented to the MT network.

For comparison, when no attentional cue was shown in a stimulation trial (Fig. 2A, left), no persistent activity was produced in the PFC network, nor was there top-down signal to modulate the MT network response during the T-period. In this example, the attended direction and the stimulus were the same ($\theta_A = \theta_S = 0°$).

As can be seen in Figure 2B, the spiking response of a neuron with $\theta_{pref} = 0°$ was enhanced by the attentional signal (red) compared with control (black). The population activity pattern (the average firing rate during the T-period plotted for all neurons) exhibits sharpened selectivity, similar to that observed experimentally (Martinez-Trujillo and Treue, 2004): neural activity was increased at the focus of attention but suppressed on the surrounds (Fig. 2C, red) compared with the unattended case (black) (selective enhancement with inhibitory surround). Such sharpening of population activity occurred because the top-down projection from PFC not only provided local excitation, but also targeted MT interneurons that then projected unspecifically onto MT excitatory neurons and suppressed firing on the flanks (supplemental Fig. 2, available at www.jneurosci.org as supplemental material). Indeed, attention strongly increased peak inhibitory firing rate (~25 Hz). When computed as a perecentage increase from baseline firing, this represents the same modulation (35%) as for excitatory neurons (Mitchell et al., 2007).

We confirmed that this enhanced selectivity was robust to parameter variations in the top-down projection, especially if changes of top-down synapses onto excitatory neurons and those onto inhibitory cells were approximately balanced (supplemental Fig. 3A,C, available at www.jneurosci.org as supplemental material). The selectivity enhancement is quantified by the ratio of firing rates with and without attention (Fig. 2C), called modulation ratio (Martinez-Trujillo and Treue, 2004), plotted as a function of the difference between the attended direction θ_A and the preference of the neuron θ_{pref}. The modulation ratio curve of the

model (Fig. 2D) was quantitatively similar to that reported experimentally (Martinez-Trujillo and Treue, 2004).

The curve can be well fitted by $G(\theta_{pref} - \theta_A) = 1.05 + 0.3\cos(\theta_{pref} - \theta_A)$. In particular, note that, for large $\theta_{pref} - \theta_A$, the modulation ratio was smaller than 1, indicating the suppression of responses away from the attentional focus (selective enhancement with inhibitory surround).

Multiplicative gain modulation of tuning curves

We next examined whether the modulation ratio of a given neuron remained constant regardless of the stimulus direction θ_S. We ran a series of simulations with a fixed attended direction while varying from trial to trial the stimulus θ_S presented during the T-period (Fig. 3A). We ob-

Figure 3. Multiplicative modulation of tuning curves. **A**, Tuning curve of an MT neuron (blue arrow in top scheme) computed from responses to the distributed range of test stimuli θ_S (black solid arrows) as attention is focused onto the preferred direction (att pref, squares) or diametrically opposed (att null, circles) or is not presented (no att, diamonds). **B**, Rescaling of tuning curves (factors, 0.73 for att pref, 1.28 for att null) reveals multiplicative scaling. **C**, Power-law relationship between the response R and the external input (sum of bottom-up I_S and top-down I_A). Data for the three cases (no att, att pref, att null) were fitted with $r = a(I_S + I_A + c)^b$ (blue curves); a and b were found to be independent of the attentional condition ($a = 1.11$ Hz/nA^b, $b = 3.97$), whereas c was smaller for attention than for nonattention conditions (0.42 and 0.59 nA, respectively).

served that the tuning curve of a neuron was multiplicatively scaled by attention from the unattended tuning curve, whether the attended direction was the preferred ("att pref," $\theta_A = 0°$) or nonpreferred ("att null," $\theta_A = 180°$) (Fig. 3B). This MGM of tuning curves was robust to modifications in the top-down input (supplemental Fig. 3B,D, available at www.jneurosci.org as supplemental material) and concurred with experimental data obtained for spatial attention (McAdams and Maunsell, 1999; Treue and Martinez-Trujillo, 1999). Several neurophysiological mechanisms for multiplicative neuronal responses have been proposed (Salinas and Abbott, 1996; Chance et al., 2002; Hansel and van Vreeswijk, 2002; Murphy and Miller, 2003), but none has been explicitly tested in a biophysical, recurrent network model constrained by multiple neurophysiological data of attentionally modulated sensory responses. We found that the relationship between the firing rate R and the total external input current (the sum of bottom-up and top-down inputs, $I_S + I_A$) in our MT excitatory cells was very well described by a power law $r = a(I_S + I_A + C)^b$ (Fig. 3C). This supports the scenario described by Hansel and van Vreeswijk (2002) and Murphy and Miller (2003), namely a power-law input–output relationship that transforms additive inputs into approximately multiplicative outputs. The fit parameters a and b are the same in all task conditions (no attention, attention to preferred motion, and attention to null motion), but the parameter c differed for attentional and nonattentional conditions. This can be readily understood as follows: c incorporates all of the rest of currents impinging on the neurons apart from I_S and I_A, and in particular it contains an important contribution from nonspecific inhibitory inputs from local circuit interneurons. Because inhibitory neurons in the MT network also receive top-down input, the inhibitory population average activity increases in attention trials and c decreases for those trials, as shown in Figure 3C. An explicit mathematical description of how this power-law input–output relationship accounts for our simulation results can be found in the supplemental Methods and supplemental Figure 5 (available at www.jneurosci.org as supplemental material).

Consistent with this multiplicative scaling of neural tuning curves, the neural firing response can be expressed as follows: $R(\theta_{pref}, \theta_A, \theta_S) = G(\theta_{pref} - \theta_A)R_0(\theta_{pref} - \theta_S)$, where $R_0(\theta_{pref} - \theta_S)$ is the neural activity in the unattended condition, and $G(\theta_{pref} - \theta_A)$ is the attention-induced multiplicative factor. Note that this equation can be used to describe both population activity

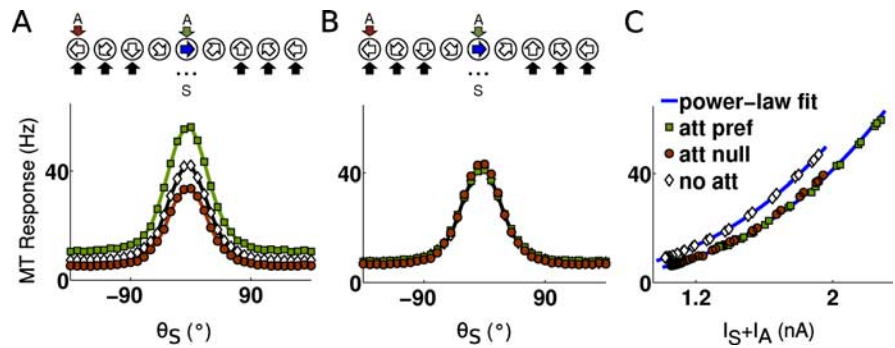(with a fixed θ_S while θ_{pref} is varied) and tuning curve of a single neuron (with a fixed θ_{pref} while varying θ_S). In particular, it can be seen from this equation why single-neuron data of Martinez-Trujillo and Treue (2004) represent population activity. In that experiment, θ_A and θ_S were covaried while keeping $\theta_A - \theta_S$ constant in all trials so that $R(\theta_{pref}, \theta_A, \theta_S)$ only depended on $\theta_{pref} - \theta_S$. Therefore, under this special condition, single-neuron responses (fixed θ_{pref}, across θ_S) are equivalent to the activity pattern of the neural population (fixed θ_S, across θ_{pref}).

Biased competition phenomenology and feature-similarity gain principle

Our model therefore reproduces salient experimentally observed effects: competition between stimuli in the sensory cortex (supplemental Fig. 1, available at www.jneurosci.org as supplemental material) (Snowden et al., 1991; Treue et al., 2000), attentional enhancement of population response selectivity (Fig. 2C) (Martinez-Trujillo and Treue, 2004), and attentional modulation of the neuronal gain (Fig. 3) (Treue and Martinez-Trujillo, 1999).

How closely does our model behavior adhere to either BC or FSGP? According to BC phenomenology, attention focused on one of two simultaneously presented stimuli should bring the activity of a neuron toward its firing rate when the attended stimulus is presented alone. FSGP, however, predicts that the attentional modulation factor of a neuron, deduced using a given stimulus, should be applicable to neural responses to other, arbitrary stimuli. We tested both of these predictions by looking at how a top-down input from the PFC network affected the processing of two stimuli presented simultaneously to the MT network during the T-period (Fig. 4A). In MT, this would correspond to two superimposed random dot patterns moving in different directions (transparent motion) (see right stimulus in Fig. 1) (Snowden et al., 1991; Treue et al., 2000), just one of them being behaviorally relevant (attended stimulus). Thus, FSGP requires that the modulation-ratio curve (Fig. 2D) should predict the population response to attended transparent motion from the corresponding activity in the unattended condition. This is indeed the case as shown in the full model simulation (Fig. 4B). The impressive agreement with the prediction demonstrates that FSGP naturally emerges in our microcircuit model. To strengthen further this point, we looked at the network response during the delay period, when no stimulus was being presented and only the baseline MT activity was affected by the top-down
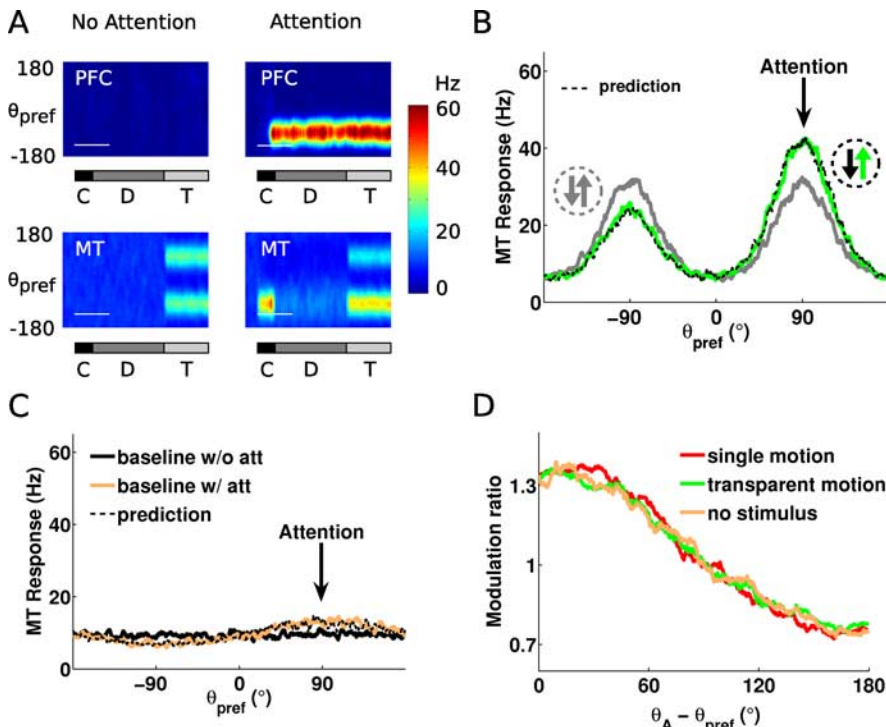
Figure 4. Attentional processing with transparent motion: the model conforms with the feature-similarity gain principle. *A*, MT population activity for trials with transparent motion presented during the test period. *B*, Test period MT activity in the attended condition (green) is well predicted by the product of the unattended activity (gray) and the modulation ratio (Fig. 2*D*). *C*, Same as *B* for baseline activity in the delay period (no sensory stimulus). *D*, Modulation ratio for single motion (red), transparent motion (green), and lack of sensory stimulation (orange).
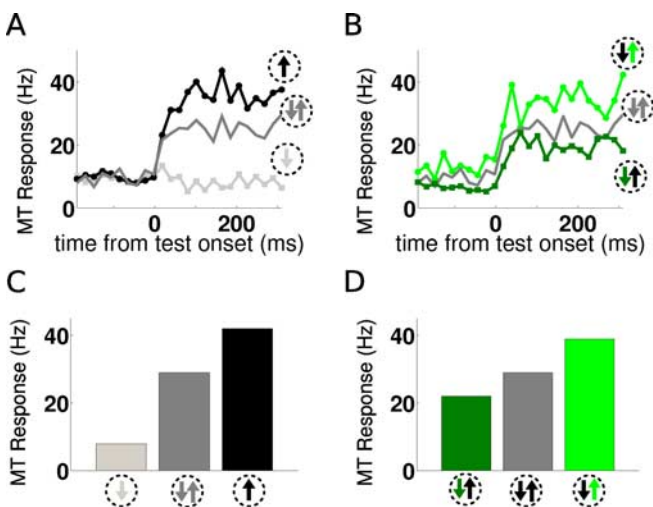

Figure 5. Neural responses to attended transparent motion: the model shows biased competition phenomenology. *A*, Dynamics of the activity of a cell at the onset of unattended test stimuli (black, preferred single motion; light gray, nonpreferred single motion; dark gray, preferred plus nonpreferred transparent motion). *B*, Same for responses to transparent motion in the unattended case (dark gray, same as in *A*) and in the attended cases (dark green, attention on preferred; light green, attention on nonpreferred). *C*, Mean activity during the test period for the cell under the conditions in *A* shows sensory competition. *D*, Same for the conditions in *B* reveals attentional bias of competitive responses.

biasing input. Indeed, there is experimental evidence for the attentional modulation of baseline activity in extrastriate visual cortex (Chelazzi et al., 1993; Luck et al., 1993; Ferrera et al., 1994; Luck et al., 1997; Kastner et al., 1999; Bisley et al., 2004). Remarkably, even for the baseline delay responses in the MT network,

which are significantly lower than evoked responses, attentional effects are extremely well predicted by the modulation ratios evaluated from the single motion stimulus presentation (Fig. 4*C*). All of this could be synthetically represented by plotting together the modulation ratio curves for each of the three types of stimulations that we used: no stimulus, single motion, and transparent motion. The three curves are indistinguishable from each other (Fig. 4*D*), confirming that our network model is a neurophysiologically plausible implementation of the FSGP (Treue and Martinez-Trujillo, 1999).

However, BC should manifest at the level of single neuron responses. As we mentioned before, we constrained the MT module network to reproduce competition in the presence of transparent motion stimuli (supplemental Fig. 1, available at www.jneurosci.org as supplemental material). As shown in Figure 5, *A* and *C* (dark gray), the response of a neuron to transparent motion, one at its preferred direction and the other at the null, is halfway between the responses when either motion component is presented alone (Fig. 5*A*,*C*, black and light gray). When attention is directed at either of the two simultaneously presented motion directions, the firing rate of neuronal response is pulled toward the condition when the attended stimulus is presented alone (Fig. 5*B*,*D*, dark and light green). Therefore, our model reproduces the BC and proves its consistency with MGM at the neuronal level (Fig. 3).

Dissecting circuit mechanisms

The neurophysiologically explicit model allowed us to dissect the circuit requirements for each of the multiple forms of attentional effects: selective enhancement with inhibitory surround of population responses, neuronal multiplicative scaling, and biased competition phenomenology. To this end, we differentially disabled certain synaptic pathways in the model. Such manipulations led to modifications in the three qualitative behaviors of the model that underlie the implementation of the attentional effects (Fig. 6*A*, data replotted from previous figures): sensory competition between stimuli (left), selective enhancement with inhibitory surround of population responses (middle), and neural gain modulation (right). When background Poisson inputs to MT neurons (see Materials and Methods) were reduced, rendering linear the input–output relationship of network neurons (Fig. 6*B*, right), the model still showed a suppressed maximal response to two stimuli with respect to isolated stimuli (competition; left panel) and attention still enhanced the selectivity of population activity by increasing activity at the peak and suppressing it at the flanks (selective enhancement with inhibitory surround; middle panel). However, neuronal multiplicative scaling no longer held as rescaled neuronal tuning curves for different attentional conditions did not overlap (right panel). Consequently, the modulation ratio depended on the stimulus presented and was different for single motion and transparent motion (right, left inset). Thus,

strong unspecific and independent background synaptic inputs were necessary in our model to satisfy the FSGP.

When the inputs to MT interneurons (from both the bottom-up afferents and MT excitatory neurons) were blocked (Fig. 6C), competition was abolished (left panel), because the response of a neuron to superimposed preferred and nonpreferred stimuli exceeded the response to the preferred stimulus alone (left panel, inset). In contrast, selective enhancement with inhibitory surround of population activity (middle panel) and multiplicative gain of tuning curves (right panel) remained intact. However, when the top-down and recurrent inputs to MT interneurons were blocked (Fig. 6D), competition (left panel) and multiplicative scaling of tuning curve (right panel) were not affected qualitatively but the selectivity of population activity was no longer sharpened through off-focus suppression, abolishing both selective enhancement with inhibitory surround (middle panel) and correct attentional biasing of stimuli competition (middle panel, inset). Blocking only one of the pathways, rather than a pair simultaneously, was not sufficient to induce the observed qualitative changes (supplemental Fig. 4, available at www.jneurosci.org as supplemental material). These manipulations showed that BC was generated through an interplay of sensory competition (Fig. 6C) and selective enhancement with inhibitory surround of population activity (Fig. 6D), whereas FSGP relied essentially on the neuronal multiplicative mechanism (Fig. 6A). Thus, the two conceptual principles, BC and FSGP, can be mechanistically dissociable.

Discussion

We present here a biophysically detailed microcircuit model of two interacting cortical areas that integrates top-down input generation in a working memory module and the ensuing attentional modulations in a sensory module. We show that such a network is able to account for most experimental evidence of feature-based attentional modulations of neural and population responses in area MT.

The biophysical detail of our model allows us to formulate the plausible mechanistic bases of current conceptual interpretations of attentional processing at the microcircuit level. In particular, we provide the first biophysical instantiation of the feature-similarity gain principle (Treue and Martinez-Trujillo, 1999) in a cortical network, which is also compatible with the biased competition principle. Notice that we proved the consistency between FSGP and BC in a model that can only represent a single feature of the stimulus. To make explicit the interactions between these two principles in the context of mixed featural and spatial attention tasks would require a significantly more complex computational
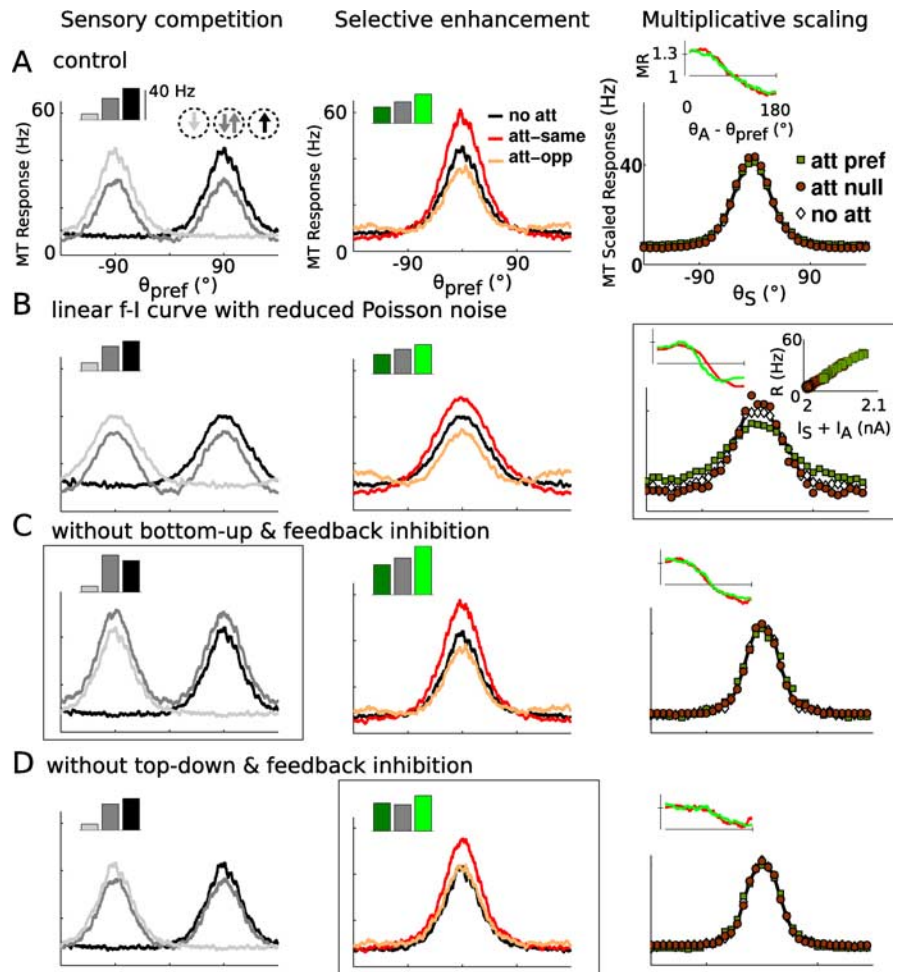

Figure 6. Mechanisms of the feature-similarity gain principle and biased competition in the model. Left, MT population activity for single and transparent motion in unattended condition (inset, activity of a single neuron; colors are as in Fig. 5C). Middle, MT population activity for single motion in attended and unattended conditions (inset, single neuron responses to transparent motion; colors are as in Fig. 5D). Right, Rescaled MT tuning curves as in Figure 3B. Modulation ratios for single (red), and transparent (green) motion in left inset. **A**, Control data (from Figs. 2–5). **B**, Reduced external Poisson input (right inset, f–I curves turn linear). Parameter modifications were $G_{EE}^{PFC \to MT} = 0.012$ nS, $G_{EI}^{PFC \to MT} = 0.005$ nS, $I_0^E = 2$ nA, $I_1^E = 0.05$ nA, $I_0^I = 0$, $I_1^I = 0$, $g_{ext,E} = 1$ nS, and $g_{ext,I} = 1$ nS. **C**, Blockade of bottom-up and local excitatory inputs to MT interneurons. Parameter modifications were $G_{EEPFC}^{\to MT} = 0.134$ nS, $I_0^E = 0.225$ nA, $I_0^I = 0$, $I_1^I = 0$, $G_{EI,AMPA}^{MT} = 0$, $G_{EI,NMDA}^{MT} = 0$, and $g_{ext,I} = 7.12$ nS. **D**, Blockade of top-down and local excitatory inputs to MT interneurons. Parameter modifications were $G_{EE}^{PFC \to MT} = 0.061$ nS, $G_{EI}^{PFC \to MT} = 0$, $I_0^E = 0.9$ nA, $I_1^E = 0.18$ nA, $G_{EI,AMPA}^{MT} = 0$, $G_{EI,NMDA}^{MT} = 0$, and $g_{ext,I} = 7.12$ nS. When not indicated, axis scales are as in corresponding control case (**A**). Framed graphs show departure from these qualitative behaviors.

model that incorporates two continuous dimensions of the stimulus.

We want to emphasize here that, although there is general agreement on the significance of biased competition phenomenology for attention, there is a diversity of views on how this might be mechanistically implemented (Desimone, 1992). Indeed, although the original formulation of the biased competition interpretation (Desimone and Duncan, 1995) did not provide a specific mechanistic framework, some later schematical mathematical implementations (Reynolds et al., 1999; Reynolds and Chelazzi, 2004) suggested that attentional modulations could alter feedforward inputs to extrastriate cortex to generate BC phenomenology. This mechanism underlies a number of computational network models of BC (Hamker, 2004, 2005). Other mechanistic models of biased competition, instead, attribute attentional modulations to local inhibitory interactions between neurons activated both by bottom-up and top-down

inputs (Deco and Rolls, 2005). Our model is also a BC model and is in line with this last family of models, and it shows how they can also manifest the feature-similarity gain principle. Note that other BC models (Reynolds et al., 1999, 2000; Reynolds and Chelazzi, 2004; Deco and Rolls, 2005) have also shown the capability to integrate consistently other neurophysiological effects of attention (biasing of baseline activity, modulation of responses to single stimuli) but not FSGP, and not with the mechanistic detail of the model reported here.

An important ingredient in our network is the normalization of inputs in area MT (supplemental Fig. 1, available at www.jneurosci.org as supplemental material). Normalization of responses to combined stimuli is a common characteristic of many visual areas (Snowden et al., 1991; Qian and Andersen, 1994; Carandini et al., 1997; Britten and Heuer, 1999).

This has been incorporated in various mathematical models, typically using some phenomenological construct such as divisive inhibition (Heeger, 1992; Carandini et al., 1997; Simoncelli and Heeger, 1998; Reynolds et al., 1999; Rust et al., 2006). In our biophysical network, normalization does not emerge from shunting inhibition (Holt and Koch, 1997) but from divergent feedforward inhibition, and it underlies the competition of stimuli in BC phenomenology (Fig. 6C).

Although our work suggests a unified and biophysically based model for feature-similarity gain principle and biased competition, we also identified dissociable circuit components that are critical to each of the two attentional effects. This is in contrast to a previous conceptual model suggesting that feature-similarity gain principle necessarily implies biased competition and vice versa (Boynton, 2005). Indeed, by selectively inactivating various pathways in our model, we found that our circuit model could implement either effect independently of the other (Fig. 6), thus proving that these conceptual models are not necessarily associated. Because experimental evidence indicates that attentional systems comply with both biased competition and with the feature-similarity gain principle, this imposes severe constraints on viable mechanistic schemes of neural circuits for attentional selection. We have presented here the first physiological model that accounts for both explicitly.

Anatomically, it is known that PFC and MT are reciprocally connected (Barbas, 1988; Schall et al., 1995; Burman et al., 2006). In a recent paper, Zaksas and Pasternak (2006) have examined neurons in MT and dorsolateral PFC in a delayed match-to-sample task using random-dot motion patterns. It was found that neurons in the dorsolateral PFC showed decaying selectivity for the cue stimulus through the delay period of the task, suggesting that they were not responsible for the memory maintenance, although they did retain information for a longer time than MT neurons during the delay period. Therefore, information storage for motion direction in working memory might reside in a frontal (or parietal) area that is different from dorsolateral PFC. Alternatively, working memory might implicate coding strategies other than sustained activation of neuronal firing (Zaksas and Pasternak, 2006), in contrast to the recurrent network model of working memory (Durstewitz et al., 2000; Wang, 2001; Constantinidis and Wang, 2004; Compte, 2006). However, the selectivity of MT neurons during the delay of a working memory task has not been detected in some studies (Ferrera et al., 1994) or has been found not to be stable over time (Bisley et al., 2004; Zaksas and Pasternak, 2006), possibly in relation to the matching-to-sample task design. In any event, the interaction between MT and PFC has not yet been analyzed physiologically in attention tasks, and will be needed to directly test our model predictions.

Our model can be extended in future research in several important ways. In particular, it has been reported that selective attention modulates sensory responses in a manner consistent with a contrast enhancement of the attended stimulus (Reynolds et al., 2000; Martinez-Trujillo and Treue, 2002) (but see Williford and Maunsell, 2006). It will be interesting to incorporate in our model contrast dependence and contrast adaptation that interacts with attentional signaling. Another recent finding that is capturing great interest is the attentional modulation of gamma band coherency of neural activity in sensory cortices (Fries et al., 2001; Bichot et al., 2005; Womelsdorf et al., 2006). Our model network does show attention-induced oscillatory activity in the gamma band, which emerges from the interplay of AMPA-type excitation and slower GABA$_A$ inhibition in the microcircuit (Compte et al., 2000). Preliminary data (data not shown) indicate that gamma oscillations do not affect the findings reported here qualitatively; but a thorough exploration of this issue will be undertaken in a separate work. Note that our model of spiking networks for both the sensory and the attentional source circuits is particularly well suited for elucidating the role of oscillation and synchronicity in attentional processing.

The consistent integration in a single physiological model of a significant amount of experimental data from MT, which is also shared by other extrastriate areas selective to different stimulus features, leads us to hypothesize that it contains the fundamental mechanistic backbone of a canonical attentional circuit. Indeed, key electrophysiological observations that are accounted for by our network model, sensory competition (Snowden et al., 1991; Treue et al., 2000), biased competition phenomenology (Treue and Maunsell, 1996), population activity selectivity enhancement (Martinez-Trujillo and Treue, 2004), and tuning curve multiplicative scaling (Treue and Martinez-Trujillo, 1999), have also been observed in other extrastriate cortical areas (notably V4), but all of these observations have not been conjointly reported in one same visual area apart from MT with a consistent attentional protocol. We propose that this basic architecture is a canonical model for attentional selection in the cortex that is replicated for each relevant attendable parametric feature of the stimulus, including orientation, spatial location, etc. (Maunsell and Treue, 2006).

Our hypothesis that the network architecture that we put forward here constitutes a canonical model that applies to all stimulus features coded by different neural systems along the visual pathway generates some specific predictions. First, the model suggests that population selectivity enhancement holds valid for other systems, such as modulation of receptive fields of neurons in V4 by spatial attention. This can be tested with neural recordings, by testing the RF of a neuron with an optimal stimulus positioned at different locations in the RF, once while the monkey passively fixates (unattended condition) and again while the monkey attends to the stimulus so that the attended location and stimulus location covary from trial to trial (attended condition) [paralleling the experiment in MT by Martinez-Trujillo and Treue (2004)]. Alternatively, the enhancement of population response selectivity could be explored directly using imaging techniques. Our model predicts that the attended response pattern will have sharpened selectivity with respect to the unattended profile, through the enhancement of activity in the center and suppression of activity in the periphery. Second, a clear distinction between selectivity enhancement for population activity and multiplicative scaling of single neural tuning curve should help disambiguate seemingly contrasting experimental results, such as previously reported selectivity modulation (Spitzer et al., 1988)

and gain change (McAdams and Maunsell, 1999) of orientation selectivity in V4 neurons. Although McAdams and Maunsell (1999) gave a reasonable explanation of this discrepancy based on the different way in which selectivity measurements were taken in either study, our model suggests that selectivity enhancement at the level of population activity could have been an additional source of confounds. Our prediction is that appropriate experiments designed specifically to compare modulations of single neuron and population activity by attention to orientation, when decoupled from spatial attention, should disambiguate the contrasting reports and reveal the neurophysiological fingerprints of our canonical attentional model: competition, population selectivity enhancement, and single neuron multiplicative scaling. Third, the feature-similarity gain principle has so far been tested only with one stimulus (Martinez-Trujillo and Treue, 2004); our model predicts that the modulation ratio thus obtained can be used to predict MT population response to other stimuli, such as transparent motion (Fig. 4*B*). This is a critical test for the feature-similarity gain principle. An analogous prediction applies to V4 neural responses when attention selects one of the orientations in a plaid pattern. More generally, it would be intriguing to see whether the attentional modulation ratio can predict neural responses when attention is directed to one of an arbitrary number of stimuli in a cluttered visual scene, in visual search tasks. Finally, our model suggests that distinct synaptic pathways are critical to different aspects of attentional modulation (Fig. 6); experimental progress in this direction would represent a major step toward a mechanistic understanding of selective attention at the microcircuit level.

References

Barbas H (1988) Anatomic organization of basoventral and mediodorsal visual recipient prefrontal regions in the rhesus monkey. J Comp Neurol 276:313–342.

Barceló F, Suwazono S, Knight RT (2000) Prefrontal modulation of visual processing in humans. Nat Neurosci 3:399–403.

Bichot NP, Rossi AF, Desimone R (2005) Parallel and serial neural mechanisms for visual search in macaque area V4. Science 308:529–534.

Bisley JW, Zaksas D, Droll JA, Pasternak T (2004) Activity of neurons in cortical area MT during a memory for motion task. J Neurophysiol 91:286–300.

Born RT, Bradley DC (2005) Structure and function of visual area MT. Annu Rev Neurosci 28:157–189.

Boynton GM (2005) Attention and visual perception. Curr Opin Neurobiol 15:465–469.

Britten KH, Heuer HW (1999) Spatial summation in the receptive fields of MT neurons. J Neurosci 19:5074–5084.

Burman KJ, Palmer SM, Gamberini M, Rosa MGP (2006) Cytoarchitectonic subdivisions of the dorsolateral frontal cortex of the marmoset monkey (*Callithrix jacchus*), and their projections to dorsal visual areas. J Comp Neurol 495:149–172.

Carandini M, Heeger DJ, Movshon JA (1997) Linearity and normalization in simple cells of the macaque primary visual cortex. J Neurosci 17:8621–8644.

Chance FS, Abbott LF, Reyes AD (2002) Gain modulation from background synaptic input. Neuron 35:773–782.

Chelazzi L, Miller EK, Duncan J, Desimone R (1993) A neural basis for visual search in inferior temporal cortex. Nature 363:345–347.

Colby CL, Goldberg ME (1999) Space and attention in parietal cortex. Annu Rev Neurosci 22:319–349.

Compte A (2006) Computational and in vitro studies of persistent activity: edging towards cellular and synaptic mechanisms of working memory. Neuroscience 139:135–151.

Compte A, Brunel N, Goldman-Rakic PS, Wang X-J (2000) Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. Cereb Cortex 10:910–923.

Constantinidis C, Wang X-J (2004) A neural circuit basis for spatial working memory. Neuroscientist 10:553–565.

Corbetta M, Shulman GL (2002) Control of goal-directed and stimulus-driven attention in the brain. Nat Rev Neurosci 3:201–215.

Deco G, Rolls ET (2005) Neurodynamics of biased competition and cooperation for attention: a model with spiking neurons. J Neurophysiol 94:295–313.

Desimone R (1992) Neural circuits for visual attention in the primate brain. In: Neural networks for vision and image processing (Carpenter GA, Grossberg S, eds), pp 343–364. Cambridge, MA: MIT.

Desimone R, Duncan J (1995) Neural mechanisms of selective visual attention. Annu Rev Neurosci 18:193–222.

Durstewitz D, Seamans JK, Sejnowski TJ (2000) Neurocomputational models of working memory. Nat Neurosci 3 [Suppl]:1184–1191.

Fan J, McCandliss BD, Fossella J, Flombaum JI, Posner MI (2005) The activation of attentional networks. NeuroImage 26:471–479.

Ferrera VP, Rudolph KK, Maunsell JH (1994) Responses of neurons in the parietal and temporal visual pathways during a motion task. J Neurosci 14:6171–6186.

Fries P, Reynolds JH, Rorie AE, Desimone R (2001) Modulation of oscillatory neuronal synchronization by selective visual attention. Science 291:1560–1563.

Grent-t Jong T, Woldorff MG (2007) Timing and sequence of brain activity in top-down control of visual-spatial attention. PLoS Biol 5:e12.

Hagler DJJ, Sereno MI (2006) Spatial maps in frontal and prefrontal cortex. NeuroImage 29:567–577.

Hamker FH (2004) A dynamic model of how feature cues guide spatial attention. Vis Res 44:501–521.

Hamker FH (2005) The reentry hypothesis: the putative interaction of the frontal eye field, ventrolateral prefrontal cortex, and areas V4, IT for attention and eye movement. Cereb Cortex 15:431–447.

Hansel D, van Vreeswijk C (2002) How noise contributes to contrast invariance of orientation tuning in cat visual cortex. J Neurosci 22:5118–5128.

Heeger DJ (1992) Normalization of cell responses in cat striate cortex. Vis Neurosci 9:181–197.

Holt GR, C. Koch (1997) Shunting inhibition does not have a divisive effect on firing rates. Neural Comput 9:1001–1013.

Hopf J-M, Boehler CN, Luck SJ, Tsotsos JK, Heinze H-J, Schoenfeld MA (2006) Direct neurophysiological evidence for spatial suppression surrounding the focus of attention in vision. Proc Natl Acad Sci USA 103:1053–1058.

Hopfinger JB, Buonocore MH, Mangun GR (2000) The neural mechanisms of top-down attentional control. Nat Neurosci 3:284–291.

Itti L, Koch C (2001) Computational modelling of visual attention. Nat Rev Neurosci 2:194–203.

Kastner S, Pinsk MA, De Weerd P, Desimone R, Ungerleider LG (1999) Increased activity in human visual cortex during directed attention in the absence of visual stimulation. Neuron 22:751–761.

Lebedev MA, Messinger A, Kralik JD, Wise SP (2004) Representation of attended versus remembered locations in prefrontal cortex. PLoS Biol 2:1919–1935.

Luck SJ, Fan S, Hillyard SA (1993) Attention-related modulation of sensory-evoked brain activity in a visual search task. J Cogn Neurosci 5:188–195.

Luck SJ, Chelazzi L, Hillyard SA, Desimone R (1997) Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. J Neurophysiol 77:24–42.

Martinez-Trujillo J, Treue S (2002) Attentional modulation strength in cortical area MT depends on stimulus contrast. Neuron 35:365–370.

Martinez-Trujillo JC, Treue S (2004) Feature-based attention increases the selectivity of population responses in primate visual cortex. Curr Biol 14:744–751.

Maunsell JH, Van Essen DC (1983) Functional properties of neurons in middle temporal visual area of the macaque monkey. I. Selectivity for stimulus direction, speed, and orientation. J Neurophysiol 49:1127–1147.

Maunsell JHR, Treue S (2006) Feature-based attention in visual cortex. Trends Neurosci 29:317–322.

McAdams CJ, Maunsell JHR (1999) Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. J Neurosci 19:431–441.

Miller EK, Cohen JD (2001) An integrative theory of prefrontal cortex function. Annu Rev Neurosci 24:167–202.

Mitchell JF, Sandberg KA, Reynolds JH (2007) Differential attention-

dependent response modulation across cell classed in Macaque visual area V4. Neuron 55:131–141.

Moore T (2006) The neurobiology of visual attention: finding sources. Curr Opin Neurobiol 16:159–165.

Moran J, Desimone R (1985) Selective attention gates visual processing in the extrastriate cortex. Science 229:782–784.

Murphy BK, Miller KD (2003) Multiplicative gain changes are induced by excitation or inhibition alone. J Neurosci 23:10040–10051.

Qian N, Andersen RA (1994) Transparent motion perception as detection of unbalanced motion signals. II. Physiology. J Neurosci 14:7367–7380.

Renart A, Brunel N, Wang X-J (2003) Mean field theory of irregularly spiking neuronal populations and working memory in recurrent cortical networks. In: Computational neuroscience: a comprehensive approach (Feng J, ed), pp 431–490. Boca Raton, FL: CRC.

Reynolds JH, Chelazzi L (2004) Attentional modulation of visual processing. Annu Rev Neurosci 27:611–647.

Reynolds JH, Chelazzi L, Desimone R (1999) Competitive mechanisms subserve attention in macaque areas V2 and V4. J Neurosci 19:1736–1753.

Reynolds JH, Pasternak T, Desimone R (2000) Attention increases sensitivity of V4 neurons. Neuron 26:703–714.

Rust NC, Mante V, Simoncelli EP, Movshon JA (2006) How MT cells analyze the motion of visual patterns. Nat Neurosci 9:1421–1431.

Salinas E, Abbott LF (1996) A model of multiplicative neural responses in parietal cortex. Proc Natl Acad Sci USA 93:11956–11961.

Schall JD, Morel A, King DJ, Bullier J (1995) Topography of visual cortex connections with frontal eye field in macaque: convergence and segregation of processing streams. J Neurosci 15:4464–4487.

Simoncelli EP, Heeger DJ (1998) A model of neuronal responses in visual area MT. Vis Res 38:743–761.

Snowden RJ, Treue S, Erickson RG, Andersen RA (1991) The response of area MT and V1 neurons to transparent motion. J Neurosci 11:2768–2785.

Snowden RJ, Treue S, Andersen RA (1992) The response of neurons in areas V1 and MT of the alert rhesus monkey to moving random dot patterns. Exp Brain Res 88:389–400.

Spitzer H, Desimone R, Moran J (1988) Increased attention enhances both behavioral and neuronal performance. Science 240:338–340.

Sturm W, Willmes K (2001) On the functional neuroanatomy of intrinsic and phasic alertness. NeuroImage 14:76–84.

Thiel C, Fink G (2007) Visual and auditory alertness: modality-specific and supramodal neural mechanisms and their modulation by nicotine. J Neurophysiol 97:2758–2768.

Treue S, Martinez-Trujillo JC (1999) Feature-based attention influences motion processing gain in macaque visual cortex. Nature 399:575–579.

Treue S, Maunsell JH (1996) Attentional modulation of visual motion processing in cortical areas MT and MST. Nature 382:539–541.

Treue S, Hol K, Rauber HJ (2000) Seeing multiple directions of motion-physiology and psychophysics. Nat Neurosci 3:270–276.

Wang XJ (2001) Synaptic reverberation underlying mnemonic persistent activity. Trends Neurosci 24:455–463.

Williford T, Maunsell JHR (2006) Effects of spatial attention on contrast response functions in macaque area V4. J Neurophysiol 96:40–54.

Womelsdorf T, Fries P, Mitra PP, Desimone R (2006) Gamma-band synchronization in visual cortex predicts speed of change detection. Nature 439:733–736.

Zaksas D, Pasternak T (2006) Directional signals in the prefrontal cortex and in area MT during a working memory for visual motion task. J Neurosci 26:11726–11742.