

Feature Review

Mechanisms of top-down attention

Farhan Baluch¹ and Laurent Itti^{1,2}¹Neuroscience Graduate Program, University of Southern California, Los Angeles, CA, USA²Department of Computer Science, University of Southern California, Los Angeles, CA, USA

Attention exhibits characteristic neural signatures in brain regions that process sensory signals. An important area of future research is to understand the nature of top-down signals that facilitate attentional guidance towards behaviorally relevant locations and features. In this review, we discuss recent studies that have made progress towards understanding: (i) the brain structures and circuits involved in attentional allocation; (ii) top-down attention pathways, particularly as elucidated by microstimulation and lesion studies; (iii) top-down modulatory influences involving subcortical structures and reward systems; (iv) plausible substrates and embodiments of top-down signals; and (v) information processing and theoretical constraints that might be helpful in guiding future experiments. Understanding top-down attention is crucial for elucidating the mechanisms by which we can filter sensory information to pay attention to the most behaviorally relevant events.

Introduction

Language is infused with idiomatic expressions that make explicit the distinction between bottom-up (BU) and top-down (TD) processes of attention. We might ask someone to ‘pay attention to the road’ while driving, which implies a voluntary choice to allocate resources to a subset of the perceptual input. Alternatively, we might remark that the orange sports car really ‘caught our attention’. In this case, the resource has been involuntarily captured rather than voluntarily allocated. The distinction is not limited to idiomatic expressions, but rather stems from disparate modes of attentional processing [1]. BU attention is deployed very rapidly and depends exclusively on the properties of a sensory stimulus. By contrast, TD attention is slower and requires more effort to engage.

In the modality of vision, the two modes (BU and TD) give rise to the psychophysical phenomenon of pop-out and set-size effects. In a typical visual search experiment, a subject is presented with a number of items on a display and is asked to find a target item within this display, such as a bar with a particular orientation, or color, or a combination of the two. Pop-out occurs when the target item is significantly distinct from the surrounding items (distractors), such as a horizontal bar among several vertical bars. This different item automatically attracts BU attention (or pops-out) rapidly and independently of the number of distractors [2,3]. By contrast, when the target item is distinguished only by taking into account the conjunction

of its features, such as color and orientation, BU cues alone cannot efficiently guide attention and TD attention must be recruited to scan the display. This gives rise to search times that increase with the number of distractors; in other words, a set-size effect is observed. In most real-life situations, the responses of the nervous system to a sensory input depend on both BU influences driven by the sensory stimulus and TD influences shaped by extra-retinal factors such as the current state and goal of the organism [4,5].

A distinction is also made between two types of TD mechanisms. The first type is intuitively associated with TD and is called the volitional TD process, which can exert its influence through acts of will. The second type is known as a mandatory TD process and it is an automatic, percept-modifying TD mechanism that is pervasive and that volition cannot completely eliminate. The latter TD process can develop through experience-dependent plasticity or during development, and includes contextual modulation

Glossary

BU influence: influence on the nervous system due to extrinsic properties of the stimuli.

Conjunction search: search task in which a subject is required to find a target item among several distractors, and the target is defined by a unique conjunction of features. In this type of search task, locating the target is more difficult because distractors share some of the features of the target and thus the target does not obviously stand or pop out.

Covert attention: attention paid to a subset of the sensory inputs through mental focusing.

Feed-forward sweep: first epoch of neural activity that travels from lower to higher visual areas on the onset of a visual stimulus via feed-forward connections.

Mandatory TD process: attentional process that influences sensory processing in an automatic and persistent manner.

Overt attention: attention paid through orienting of sensory organs toward a sensory input of interest.

Percept: mental impression of an external stimulus.

Pop-out search: search task in which a subject is required to find a target item among several distractors, and the target is defined by a unique visual feature not shared with any of the distractors. The target thus stands or pops out and is easy to find.

Priority map: map of visual space constructed from a combination of properties of the external stimuli, and intrinsic expectations, knowledge and current behavioral goals.

Recurrent epoch: second epoch of neural activity that occurs after an initial response to onset of a stimulus and is mediated by intra-cortical horizontal connections and inter-cortical feedback connections.

Saliency map: map of stimulus conspicuity over visual space.

Set-size effect: in search tasks, a set-size effect is observed if the time required to find the target depends on the total number of items in the display (the set size).

Task-relevance map: map of behaviorally relevant locations over visual space.

TD influence: influence on the nervous system due to extra-retinal effects such as intrinsic expectations, knowledge and goals.

Volitional TD process: attentional process that exerts influence on sensory processing through an act of volition, such as willfully shifting attention to the right part of space.

Corresponding author: Itti, L. (itti@usc.edu).

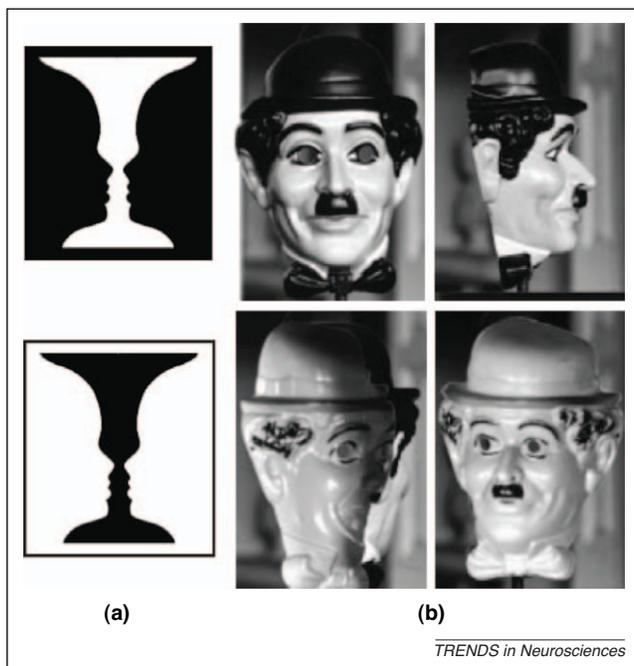


Figure 1. Mandatory versus volitional TD processes. (a) Rubin's vase illusion is an example of a volitional TD process. The percept can be switched from face to vase and vice versa through act of will, demonstrating TD modulation of perceptual processing in a volitional and dynamic manner. (b) Four frames from a demonstration of a rotating mask that seems to be convex, even in places where it is in fact concave. At the start of the rotation (top two frames), the mask is convex and, as it rotates, the viewer begins to see the inside of the mask (bottom two frames), but this still seems to be convex. This demonstrates an inherent bias in perceiving faces as convex rather than concave, even when this contradicts BU sensory information, and thus provides an example of mandatory TD processing. Reproduced with permission from [134].

of perception [5,6]. A striking example of the dichotomy between these two mechanisms is presented in Figure 1.

Previous work has extensively studied the effects of TD attention on target brain regions, including modulatory effects in early sensory areas [5,7,8]. Significant progress has been made in isolating the possible sources of TD signals [9], especially within the now well-studied fronto-parietal attention network [10]. Much less understood at present are the exact pathways, contents, meaning and form of the signals that are sent from the top down. Here, we review recent findings from physiology, lesion and computational studies that have attempted to elucidate the mechanisms and signals involved in TD modulation of sensory processing. To focus this review, we mainly concern ourselves with visual perception and the volitional TD process, although similar principles can apply in other modalities.

Brain structures and circuits of visual attention

Visual processing begins in the retina, which sends parallel streams of information to the brain through its diverse set of retinal ganglion cells and their unique interactions within the retinal circuitry [11]. A majority of retinal projections reach the lateral geniculate nucleus (LGN) and a much smaller number (approx. 10%) connect to the superior colliculus (SC). The LGN sends projections to the primary visual cortex (V1), the initial site of processing in the cortical feed-forward visual pathway. This pathway has been functionally divided into the dorsal

and ventral streams [12]. The dorsal stream has been described as the 'where' pathway and leads from area V1 to motion processing areas [medial temporal (MT) and medial superior temporal (MST)] and parietal cortices. The ventral or 'what' pathway comprises striate (V1) and extrastriate areas (V2, V3, V4) and leads to the inferotemporal cortex (IT), believed to be the last feature-selective area in the visual processing hierarchy.

Modulatory effects of attention have been observed in the constituent structures of both the dorsal and ventral streams. The first structure subject to strong attention effects is the SC. The SC is a layered midbrain structure that receives direct input from the retina, as well as feedback inputs from area V1. Salient visual events are represented in the superficial layers of the SC [13,14] and can further combine, in the deeper layers, with TD information to give rise to a priority map that guides attention [14]. This attention map is probably shared or jointly computed with the lateral intraparietal (LIP) region of the cortex [15], the frontal eye fields (FEF) [16] and visual cortices, through direct afferent connections from the cortex to the SC, as well as indirect efferent connections from the SC to the cortex via the pulvinar [17]. These connections are important for communicating attention-related signals to higher cortical areas while bypassing the canonical ventral pathway.

Situated a level above the SC in the visual processing hierarchy, are the thalamic nuclei, which are involved in processing many types of sensory information and are susceptible to modulation by attention. The LGN is the most visually responsive of the thalamic nuclei, and both physiological studies in monkeys and imaging studies in humans have shown that attention can modulate signals in the LGN [18,19]. The modulation includes enhancement of neural responses to attended stimuli and suppression of unattended stimuli [19]. Thus, visual sensory information is already subject to attentional modulation even before entering the cortex.

The first cortical stage of visual processing, area V1, is the first major feature-sensitive area of processing and is also modulated by attention. However, these effects are relatively weak [20,21]. Moving up the visual processing hierarchy from V1, V2, V4 to IT, receptive field sizes increase and visual areas are progressively more sensitive to features than spatial locations of stimuli. When attention is allocated to a certain part of visual space, neurons encoding this part are facilitated (a phenomenon known as spatial attention). The allocation of attention to a particular non-spatial feature, such as the color or orientation of an object, facilitates neurons encoding the attended feature (feature-based attention). Along the ventral pathway, extrastriate areas V4 and IT have large receptive fields and effects of feature-based attentional modulation are more evident. Motion-sensitive MT and MST areas are also modulated by both spatial and feature-based attention [22]. This tendency for combined modulation of sensory signals by both spatial and feature-based attention increases as the signals progress from lower to higher cortical areas such as the LIP.

The LIP area has been studied extensively and several excellent recent reviews have described its diverse roles in

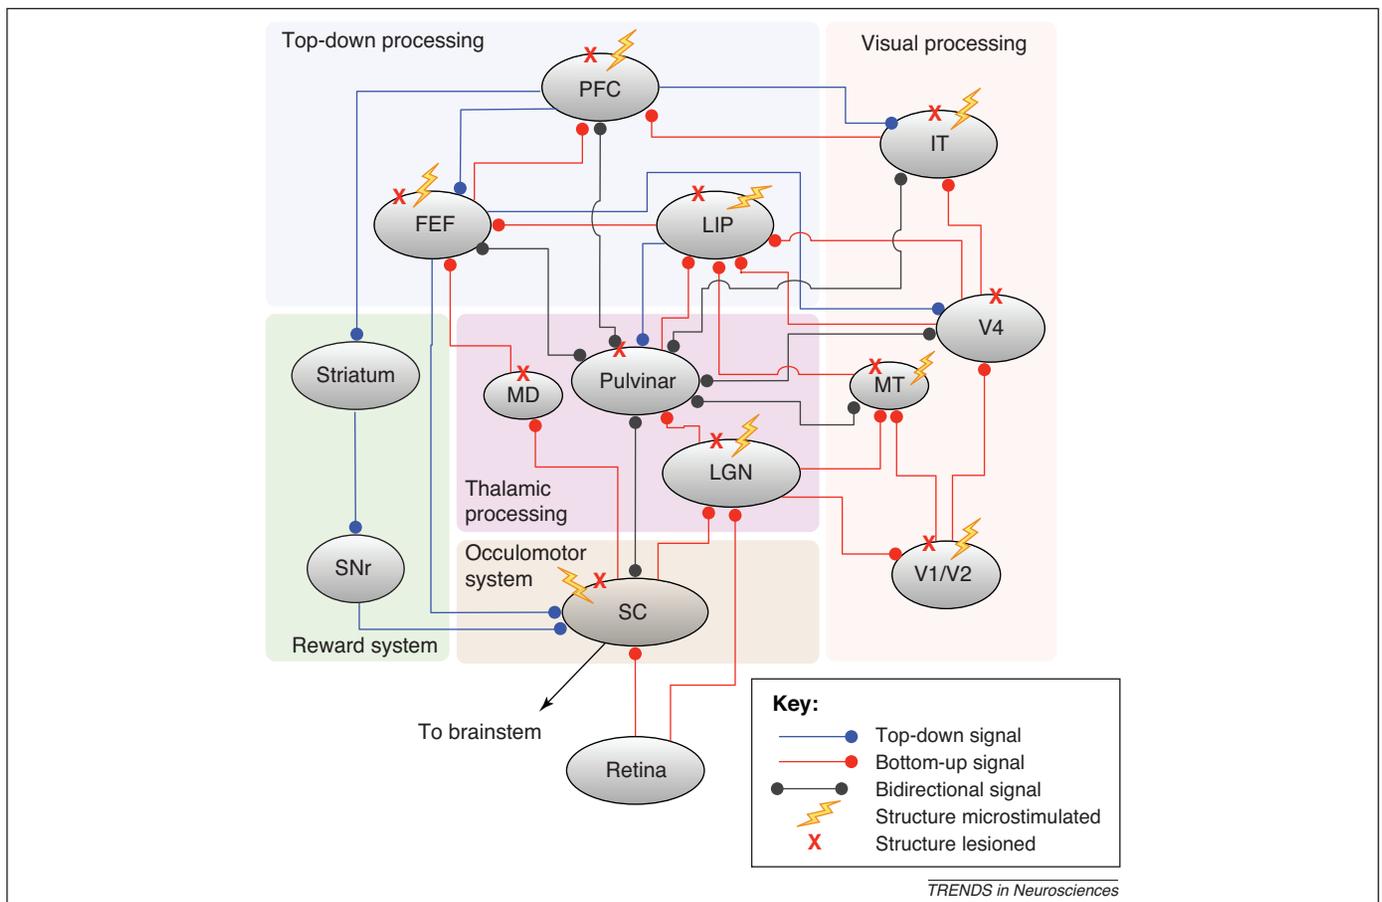
attention, reward, and oculomotor behavior [15,23,24]. It is important to point out that responses in area LIP can be driven by both BU factors, such as stimulus salience, and TD factors, such as behavioral relevance of stimuli [25], the current locus of attention [26] and oculomotor planning [15]. Therefore, the LIP is another candidate structure (beyond the SC described above) where BU and TD influences can combine to give rise to a spatial priority map [15]. The many facets of observed responses in the LIP can be attributed to the fact that both BU and a diverse set of TD influences can give rise to behavioral priority, and thus modulate LIP responses, which suggests that the LIP encodes priority in a manner largely agnostic to the factors that caused the priority [15]. Through direct feedback connections [27] or connections via the pulvinar to visual areas (see below), the LIP can communicate the fused signals to other brain areas for biasing or further attentional processing.

FEF neurons also represent salient stimuli, specifically stimuli that vary significantly from surrounding items in a visual display (known as odd-ball stimuli). The FEF has also been described as a region with neural responses characteristic of a priority map [16]. Single-unit responses in monkey FEF exhibit transients on stimulus onset, followed by a later response (latency of ~100 ms) that discriminates an odd-ball stimulus from surrounding dis-

tractors [14,16]. This suggests that the FEF computes salience in the recurrent epoch rather than the initial feed-forward sweep [28,29]. The FEF's connections to motor neurons in intermediate and deep layers of the SC make it an important structure in oculomotor behaviors associated with attention. In addition to this role of the FEF in representing BU salience, we examine in the following section its involvement in projecting TD signals to other regions of the attentional network.

Effects of attention have also been observed in prefrontal cortex (PFC). The PFC is thought to be involved in short-term memory processes, and recent studies suggest that the PFC also exhibits strong attentional selection related signals [30,31]. Owing to its involvement in short-term memory and its position high in the visual hierarchy, it is also the primary candidate for generating TD signals and sending them to sensory cortex for spatial or feature-based attentional biasing.

Therefore, the LGN, the striate and extrastriate cortex (areas V1, V4, IT and MT), as well as the SC, pulvinar, LIP, FEF and PFC, are known to be involved in attentional processes. Modulatory attentional signals are found as early as in the SC (a brainstem structure) and in the LGN, the first stop along the visual processing hierarchy [18,19]. These signals act progressively sooner and with stronger modulatory power going up from area V1 to area



TRENDS in Neurosciences

Figure 2. Flow of attentional signals in brain structures that have been implicated as being involved in attentional studies. The flash symbol indicates that a candidate structure has been microstimulated and an X indicates that the structure has been lesioned in a previous study (see Table 1 for details). The connections show the most likely type of signal being transmitted between two areas; TD signals are shown in blue, BU signals in red and bidirectional signals in gray. Abbreviations: SC, superior colliculus; SNr, substantia nigra pars reticulata; MD, mediodorsal thalamus; LGN, lateral geniculate nucleus; IT, inferotemporal cortex; MT, middle temporal area; LIP, lateral intraparietal area; FEF, frontal eye fields; PFC, prefrontal cortex.

IT [20]. These signals can bias attention for particular visual locations [32], visual features [33–36], or both. The characteristic signature of these attentional modulations onto target sensory areas includes heightened gain, sharpened tuning and other end-effects, as reviewed previously [8,37,38]. In the following section, we examine the areas that are specifically involved in mediating TD attentional signals.

Pathways of TD attention

In this section, we focus on lesion and electrophysiological studies, particularly those using methods of microstimulation and simultaneous recordings in the brain areas identified in the previous section. These areas form an attentional network (Figure 2) and we consider how TD information is relayed in this network. Microstimulation, together with reversible inactivation [using either pharmacological agents such as muscimol or transcranial magnetic stimulation (TMS)] and permanent lesion studies, have enabled researchers to go from correlation to causation in the study of perception and attention (Table 1).

It has been suggested that all sensory stimuli compete for entry into working memory [39]. Working memory not only stores information, but also enhances this information and actively generates TD attentional signals that bias feature-sensitive brain regions, and is thus vital for accomplishing behavioral goals [39]. An elegant study demonstrated that the PFC transmits the contents of working memory to the visual system by using a posterior-split-brain paradigm [40]. In this study, monkeys were presented with a visual cue in either the left or right hemifield, followed by a probe stimulus. The task was to respond to the appearance of the probe that had previously been associated with the cued item. BU signals were recorded by presenting the cue in the hemifield ipsilateral to the recording site in the IT (i.e. direct BU path from the retina up to IT), whereas TD signals could be recorded from area IT by presenting the cue in the visual hemifield contralateral to the recording site in the IT. The posterior callosum transection precluded direct communication between visual cortices from both sides of the brain, so it was hypothesized that the TD signals were fed back from the PFC to area IT (Figure 3a). To move to a more causal explanation, the next experiment involved transection of the anterior corpus callosum (thereby cutting that hypothetical pathway), which resulted in a lack of response from the IT cells [40]. These results demonstrated that TD signals correlating with working memory emanate from the PFC and feed back into the ventral stream. A more recent study also used the posterior-split-brain paradigm in conjunction with unilateral PFC removal and demonstrated that performance on a search task was mainly impaired when the goal of the search was switched on a regular basis [41]. This study thus highlighted the importance of the PFC in switching the TD context. It has also been found that microstimulation of the PFC leads to biases in target selection towards or away from the stimulation field, which demonstrates how TD signals can affect oculomotor behavior [42]. Furthermore, the sheer connectedness of the PFC suggests that its effects are pervasive and are driven by a combination of goals, rewards, salience, and planning of motor actions [9,39].

The next area proximal to the PFC, and an important player in TD attention, is the FEF. Sub-threshold FEF stimulation enhances responses of V4 neurons in the presence of a stimulus in their receptive field (Figure 4a) [43]. This demonstrates that descending TD signals from the FEF bias processing in area V4. These results were replicated in analogous regions of the barn owl [44]. The comparison of local field potentials (LFP, which may be strongly driven by afferent inputs from other brain regions) and spiking activity in the FEF (which represents intrinsic activity of FEF neurons) revealed that target-selective signals appeared in spiking activity before showing a difference in the LFP, which suggests that spatial selection was computed locally in the FEF [29]. There is speculation that this emergence of selection is communicated down to ventral regions through a synchronization of gamma-band activity between the FEF and area V4 [45]. However, a lesion study demonstrated that temporary inactivation of the FEF (using a GABA-A receptor agonist, muscimol) led to deficits not only in visually guided saccades, but also in shifts of attention during either pop-out or conjunction visual searches [46]. Contrary to an earlier study [29], these findings suggested that the FEF, although involved in covert attention, does not locally compute the selection but is rather a participant in a network with heavy involvement of the LIP.

Area LIP is strongly connected to the FEF and is integral to the attentional network through both anatomical and functional characterization. Suprathreshold microstimulation in the posterior parietal cortex (PPC), which includes both area LIP and the ventral intraparietal area (VIP), induces saccades; however, the current required to induce saccades is significantly higher compared to that required when microstimulating the FEF, which suggests that the connection from the PPC to the oculomotor system might not be a direct one. Subthreshold stimulation results in a shift of covert attention [47]. Interestingly, a non-spatial effect was also found whereby reaction times in detecting a target decreased irrespective of whether a valid, invalid, or no cue was presented [47]. This suggested that microstimulation of the LIP can override the cue signal and orient attention to the visual location corresponding to the site of stimulation. Evidence from lesion studies demonstrates that damage or inactivation of the LIP causes deficits only in the presence of multiple stimuli [48,49]. These results point to an additional role of the LIP in resolving competition among stimuli represented at lower levels through TD connections to these levels [7,50].

The aforementioned studies did not, however, differentiate between the dorsal (LIPd) and ventral (LIPv) subdivisions of the LIP. In a more recent study, the effects of local reversible inactivation (using a GABA-A receptor agonist) in areas LIPd and LIPv have been studied separately [51]. Interestingly, the many dimensions of LIP responses demonstrated previously [23] were shown to reside in disparate subdivisions of the LIP. Inactivation of the LIPd affected performance on simple saccade tasks but left visual search intact, whereas temporary lesions of the LIPv led to deficits in both search and saccadic performance [51]. The authors stressed that deficits in saccadic performance after LIPd inactivation were far smaller than

Table 1. Microstimulation and lesion studies of different brain structures involved in attention.^{a,b}

Brain region	Microstimulation studies	Refs	Lesion studies	Refs
	<i>Implications for attentional processing</i>		<i>Implications for attentional processing</i>	
SC	Shift of spatial attention	[55]	Deficit in target selection	[58,68]
	Perceptual facilitation at site of stimulation	[56]	Deficit in perceptual decision in presence of distractors	[60]
	Selection of target independent of motor plan	[57]		
	Signal transmitted to MT via Pulvinar	[70]		
LGN	Elicits visual percepts	[111]	Eliminates residual visual responses in extrastriate cortex after V1 lesion	[112]
			Disruption of smooth pursuit eye movements	[113]
Pulvinar			Deficits in target detection (human)	[110]
			No deficit in saccadic behavior	[67]
			No deficit in visual search	[68]
			Deficit in suppression of distractors during search (human)	[65]
V1			Spatial and temporal attention deficits with anterior and posterior lesions respectively (human)	[66]
	Target selection disrupted with upper layer stimulation, facilitated with lower layer stimulation	[114]	Deficit in motion detection and discrimination	[117]
	Lower current thresholds needed for evoking saccades in lower layers	[115]	Deficit in saccade targeting	[118]
V4			Median current of 5.2 μ A (6.6 μ A) required for behavioral detection of stimulation ^c	
IT/TE			Deficit in distractor suppression when target and distractor are inside RF of neuron	[119]
			Deficit in distractor suppression	[133]
	Biases perceptual judgement in visual classification	[53]	No behavioral deficit when lesion is made in infantile monkeys	[120]
MT	Bias in selection of stimulus category	[54]	Deficit in distractor suppression	[119]
	Median current of 10.3 μ A (11.3 μ A) required for behavioral detection of stimulation ^c	[116]		
LIP	Bias in motion direction discrimination	[122]	Loss in perception of motion	[124]
	Bias in motion direction during stimulus presentation but not during memorizing period	[123]	Loss in perception of motion more evident in noisy conditions	[132]
	Median current of 10.1 μ A required for behavioral detection of stimulation ^c	[116]		
FEF	Sub and suprathreshold stimulation lead to covert and over shifts of attention respectively	[47]	Deficit in distractor suppression even when stimuli are non overlapping within RF, contrast with [121]	[49]
	Bias in visual selection	[125]	Dorsal lesion leads to oculomotor deficits ventral lesion leads to attentional and oculomotor deficit	[51]
			Affects performance in tasks requiring spatial attention	[48]
PFC	Enhanced response elicited in V4	[43]	Deficit in target detection	[46]
	Facilitation akin to allocation of covert attention	[126]	Enhanced contrast sensitivity in fovea but not periphery (human)	[128]
	Bias toward direction of saccade plan rather than location of attention	[127]	Disruption of facilitation by saccade plan to location corresponding with stimulation site (human)	[129]
PFC	Bias in target selection	[42]	Loss of TD signal recorded in IT	[40]
	Disruption in saccadic activity	[130]	Decrease in behavioral performance when cue is frequently switched	[41]
			Elimination of acetylcholine release in sensory cortex after stimulus presentation (rat)	[131]

^aAll studies have been conducted in monkeys unless otherwise denoted.

^bRF, receptive field; SC, superior colliculus; LGN, lateral geniculate nucleus; IT, inferotemporal cortex; MT, middle temporal area; LIP, lateral intraparietal area; FEF, frontal eye fields; PFC, prefrontal cortex.

^cStimulation current values reported in two monkeys (see [116] for details).

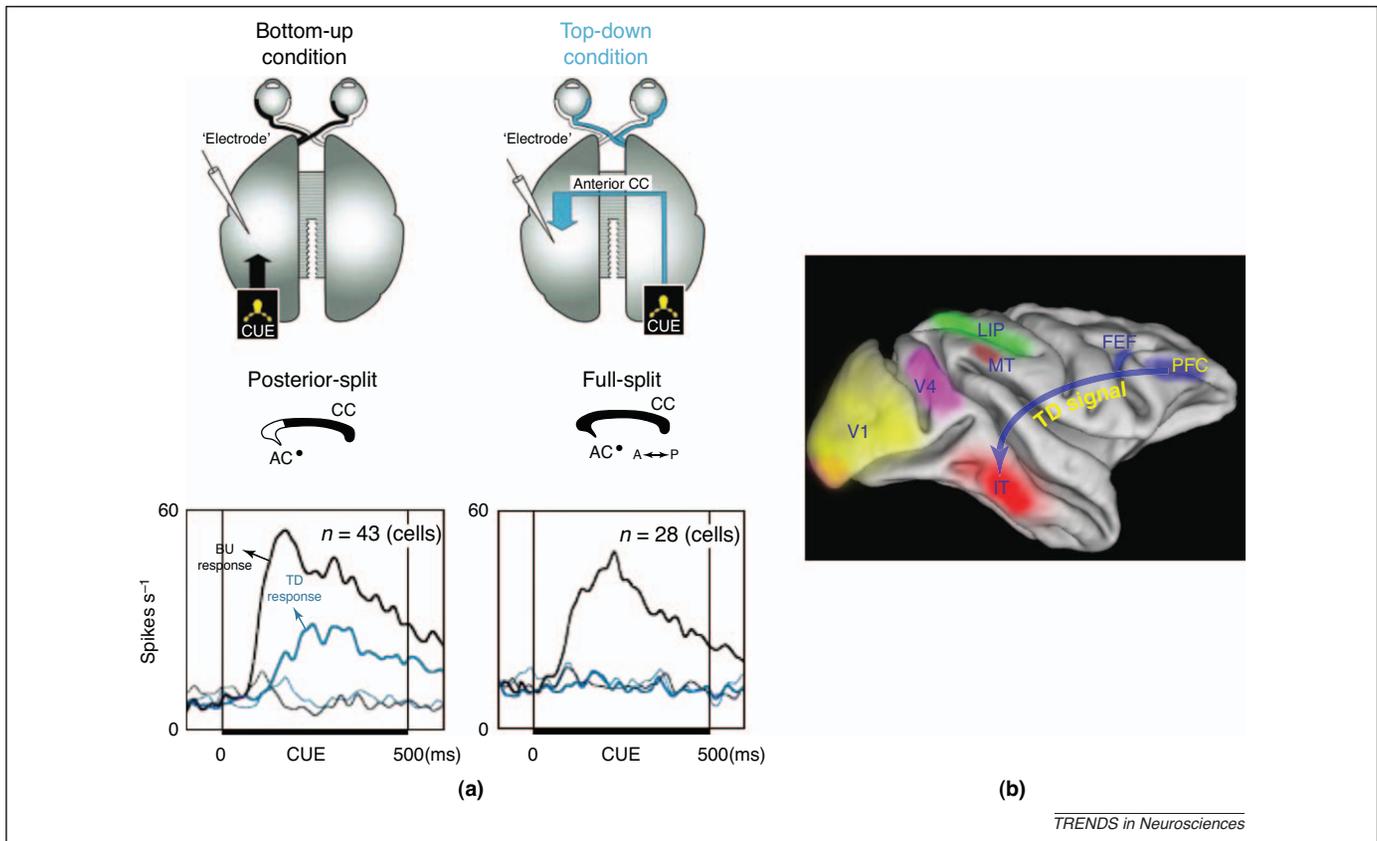


Figure 3. Role of the PFC in mediating TD attentional signals. **(a)** Posterior-split-brain paradigm. In this study, monkeys had to associate stimulus A with stimulus B [40]. Stimulus A was then presented as the cue, followed by a probe stimulus, and the task was to release a lever when the probe matched the associated stimulus B. Neurons in the inferotemporal cortex (IT) were recorded in one hemisphere while the cue was presented either contralateral to the recording site (BU condition, i.e. information about the cue could reach area IT directly; top-left panel) or ipsilateral (TD condition, i.e. information about the cue could only reach area IT via the anterior corpus callosum; top-right panel). The bottom panels show neural responses in BU (black trace) and TD (blue trace) conditions. The left-hand plot shows the responses after a posterior split, demonstrating how cue information could reach area IT in both the BU and TD conditions. The right-hand plot shows complete abolition of the TD signal after a full split of the corpus callosum (CC). This is one of the clearest demonstrations of the two different types of signal, BU and TD, recorded in a visual area. Reproduced with permission from [40]. **(b)** 3D rendering of macaque monkey brain showing regions involved in visual processing and TD attention. The areas include the first visual area (V1), fourth visual area (V4), medial temporal cortex (MT), lateral intraparietal cortex (LIP), frontal eye fields (FEF), inferotemporal cortex (IT) and prefrontal cortex (PFC). The blue arrow shows the pathway for TD signals investigated in the experiment shown in (a). Rendering of the brain was done using a macaque atlas data set [135] processed using the Caret software [136].

those observed after inactivation of the FEF and SC. They thus concluded that the LIP might influence or modulate the motor decision but that the final decision is made by more downstream structures such as the SC and FEF. This coincides with the view that attentional selection might indeed be separate from motor selection [23]. As discussed previously, the diversity of properties exhibited by LIP neurons might reflect the fact that it encodes priority without regard for what caused the priority, BU or TD influences.

We now consider feature-selective visual areas V1, V2, V4, MT and IT. These visual processing areas drive BU attentional signals and are targets for TD attentional biasing signals. For accurate biasing of sensory signals, specific local circuitry and the nature and size of receptive fields in each of these areas must constrain the nature and granularity of TD signals. Two types of feedback signals from higher cortical regions or thalamus can influence the visual processing areas [52]. One type of feedback signal can flow between a higher visual processing area to a lower one within the visual processing hierarchy (Figure 4b). Another type of feedback signal can flow between an

attention area such as the FEF and a processing area such as area V4. Figure 4a presents data from a study that demonstrates a specific example of this type of feedback signal [43]. The flow of TD attentional signals from the PFC to area IT is another example of how TD attentional signals from higher cortex can influence a feature sensitive sensory area [40]. Microstimulation in area IT results in biases of object recognition [53], or even of face detection when microstimulating face-selective sites within area IT [54]. The striate and extrastriate cortices, therefore, are all amenable to modulation by TD attention through feedback connections from higher to lower visual areas.

In summary, TD signals can emerge from the PFC to bias visual cortices through direct connections, such as from the PFC to area IT, or possibly through the pulvinar (see below). Similarly, there is evidence that a direct connection from the FEF to area V4 might exist, which further demonstrates the possible communication of TD information from higher cortex to sensory areas. TD signals from the PFC probably contain detailed information about the target and this information might be used to bias feature-selective areas of sensory cortex. The FEF and LIP,

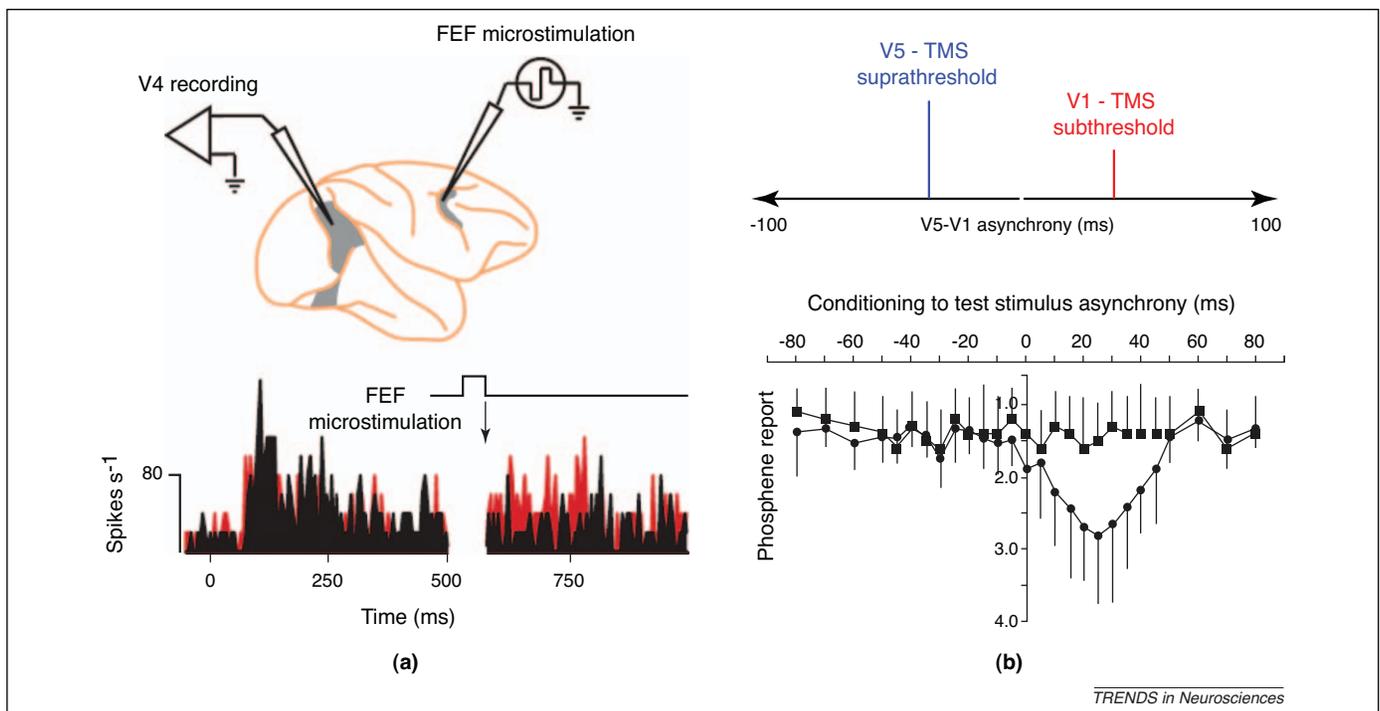


Figure 4. Role of feedback from higher to lower cortical areas in mediating attention and perception. **(a)** Neuronal activity from visual area V4 was recorded in monkeys simultaneously as the frontal eye field (FEF) was microstimulated (top panel). Histogram of neuronal activity in area V4 (bottom panel) in the control condition (black) and the stimulation condition (red). Clear enhancement of the response is evident after FEF stimulation. This demonstrates the role of frontal areas in modulating responses in a sensory visual area such as area V4. Reproduced with permission from [43]. **(b)** Visual area V5 in human subjects was stimulated with suprathreshold transcranial magnetic stimulation (TMS) pulses, followed by subthreshold TMS stimulation of visual area V1 [137]. The top panel shows the TMS paradigm used. The bottom panel shows a plot of subjective report by human subjects of phosphene perception resulting from TMS stimulation as a function of time lag between V1 and V5 stimulations (negative X values correspond to area V1 stimulation before area V5). A Y-value of 1.0 indicates that the subject perceived that a phosphene was present and moving; a value of 2.0 indicates that a phosphene was present but the subject was uncertain of motion; and a value of 3.0 indicates that the subject could see the phosphene but it was stationary. Results show that disruption of V1 activity between 5 and 45 ms after V5 stimulation results in the absence of motion, which thus demonstrates the importance of feedback signals to early visual areas for the perception of motion. Reproduced with permission from [137].

in particular, might host spatial maps encoding the behavioral relevance of visual space dependent on both BU and TD factors.

Subcortical influences on TD attention

Evidence suggesting that cortical areas have a strong influence on attention was discussed in the previous section. There are also several subcortical areas that play a crucial role in defining and communicating attentional signals (Figure 5a). It has been demonstrated that the phenomenon of change blindness, in which changes to a particular part of a visual scene go undetected, could be eliminated in monkeys by placing an attention-grabbing salient stimulus in the location where the blindness occurs [55]. Interestingly, the same effects were also observed by microstimulating the SC where receptive fields overlapped with the region of blindness [55]. This demonstrated that stimulation of the SC is equivalent to adding salience to a region of space; in other words, the SC can strongly bias attentional deployment. Another study demonstrated enhanced behavioral performance on a perceptual task with stimuli at locations corresponding to the site of stimulation in the SC [56], mimicking the effects of a shift of attention.

In another study, microstimulation of the SC in monkeys led to a bias in target selection decisions [57], which demonstrates that the SC is also involved in target selection. Conversely, inactivation of the SC led to target selection errors [58]. The SC is therefore involved in both

attentional selection and saccadic behavior. One study was able to elegantly dissociate saccade preparation signals from attentional signals [59], which clarified any ambiguity about the dual roles of the SC in oculomotor behavior and attentional control. This study involved recording from visual, visuomotor, and motor neurons in monkey SC. This revealed that visuomotor neurons encode the shift of covert attention (Figure 5b). It has also been shown that the SC is involved in gating covert attention signals used for making perceptual decisions by higher cortical areas [60].

The SC connections to the FEF and LIP, together with its role as an oculomotor structure, make it an important structure in mediating covert and overt attention. Furthermore, given its direct involvement in oculomotor behavior, it has been suggested that the SC could host the final priority map that guides attention based on a fusion of TD and BU attentional signals received from cortex and elsewhere [14].

Moving up the neuraxis to the thalamus, three important nuclei associated with visual functions are found: the LGN, the thalamic reticular nucleus (TRN) and the pulvinar nucleus. The LGN and TRN modulate their signals in a reciprocal manner (Figure 5c). When monkeys attended inside the receptive field of the recorded TRN neuron, the responses of this cell were reduced, whereas responses in the LGN were enhanced [18]. This reciprocal response in the TRN and LGN neurons was found in the initial phase of

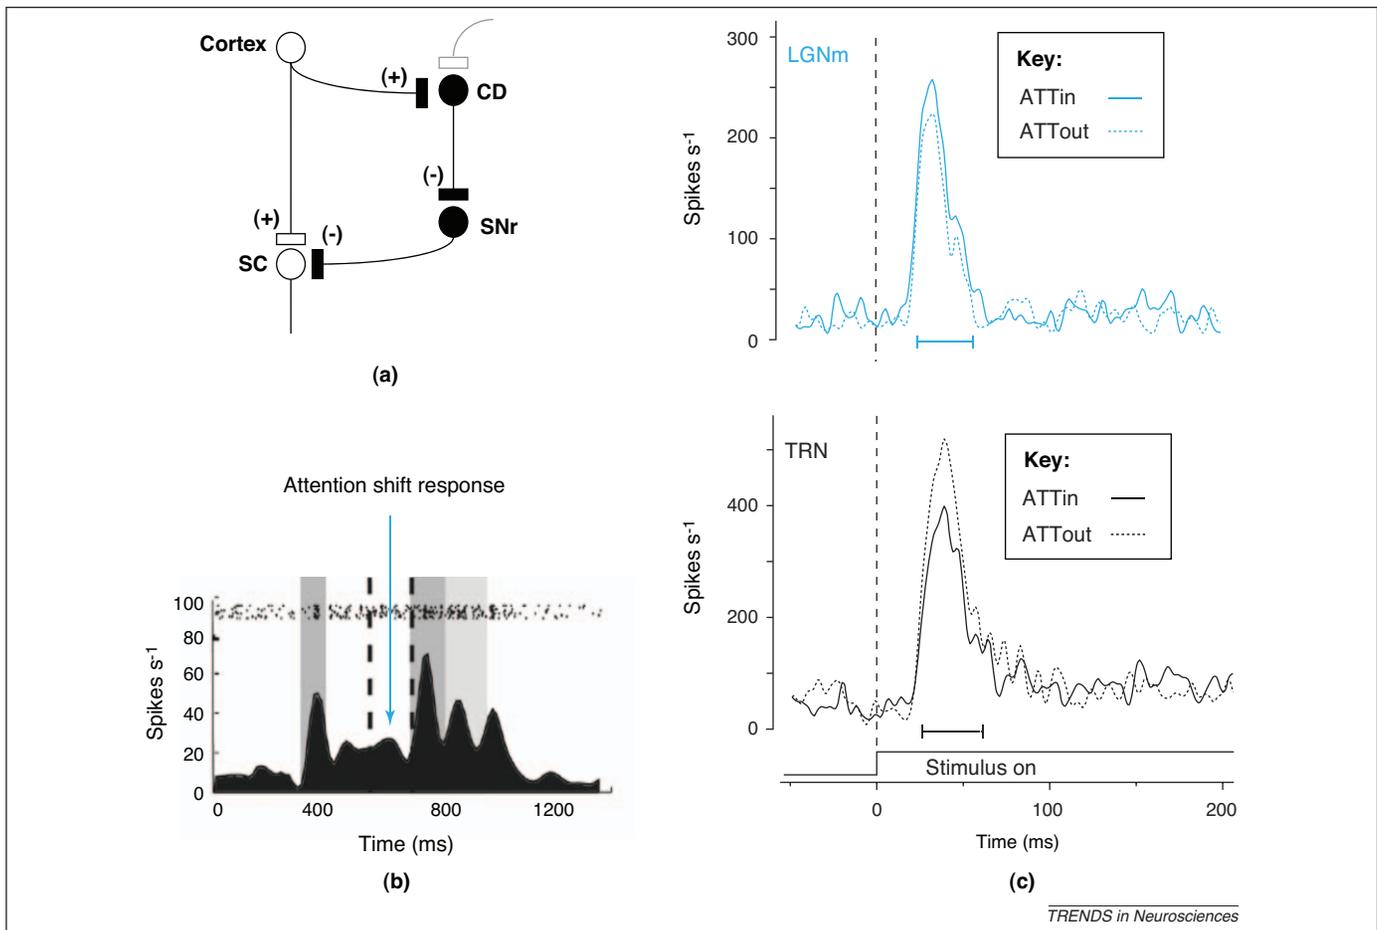


Figure 5. Role of subcortical structures in attention. **(a)** Schematic drawing of circuitry that has been proposed to be involved in the generation of eye movements towards locations of reward [78]. The cortex sends excitatory inputs to both the superior colliculus (SC) and the caudate nucleus (CD). The CD in turn inhibits the substantia nigra pars reticulata (SNr), which then reduces its tonic inhibition on the SC. A disinhibited SC enables eye movements to be made. Reproduced with permission from [78]. **(b)** Neuronal activity from a visuomotor cell in the SC. Monkeys were first presented with a spatial cue, followed by an oriented stimulus at the cued location. Monkeys then made a saccade in the direction corresponding to the orientation of the stimulus. The orientation was always orthogonal to the location of the cue, and this dissociates shifts of attention from saccadic behavior. The plot shows responses of a visuomotor SC cell, which shows significant activity in the attention shift period (between the dashed lines) that occurs immediately after presentation of the cue, whereas purely motor cells in the deeper layers of the SC did not show such a response (data not shown). Reproduced with permission from [59]. **(c)** Neuronal activity recorded from the thalamus in awake behaving monkeys. The monkeys were presented with a central cue that instructed them to attend to one of two peripheral oriented bar stimuli, one inside the receptive field (RF) of a recorded neuron and one outside the RF. The top shows the spike density of a magnocellular lateral geniculate nucleus (LGNm) neuron that exhibits an enhanced response when the monkey attends to a stimulus inside the RF (ATTin condition) of the neuron compared to when the monkey attends to a stimulus outside the RF (ATTout condition). The bottom shows responses in the thalamic reticular nucleus (TRN), which responds in a reciprocal manner to the LGNm neuron exhibiting an enhanced response when attention is allocated to a stimulus outside the RF. Therefore, the TRN might gate responses in the LGN. Reproduced with permission from [18].

the response to a visual stimulus. In a later phase, the TRN response remained unchanged, but attention further enhanced responses in LGN. These results suggest that (i) the TRN serves as the initiator of modulation in the LGN and (ii) attentional modulation begins at an early stage in the LGN. The TRN therefore plays a crucial role in modulating visual signals at a very early stage of processing.

The pulvinar is a hyperconnected nucleus of the thalamus that has been implicated in the function of visual attention based on anatomical [17,61–63], physiological [64], lesion [65–68] and computational [69] studies. It has been shown that a monkey's ability to suppress distractors is diminished when the pulvinar is pharmacologically inactivated via administration of muscimol [64]. Relay neurons have also been identified in the pulvinar by microstimulating the SC and area MT while simultaneously recording from cells in the pulvinar [70]. This study adds to evidence of a subcortical route for visual signals to reach higher cortex via the pulvinar. At the same

time, its bidirectional connections with higher cortical areas make it a potentially important structure in mediating TD signals. However, the pulvinar remains an understudied nucleus, and further studies on this particular brain nucleus are warranted.

Subcortical structures, therefore, both modulate signals in areas encoding BU and TD information, such as the LIP and FEF, and receive TD information from higher cortical areas, directly or possibly through the pulvinar. The SC itself is believed to host a priority map, but this priority map might have closer correspondence to representations needed for motor decisions, including oculomotor behavior and head movements. Thalamic nuclei, including the LGN and TRN, modulate visual signals early on, before they reach cortex, and the pulvinar might be a key relay in communicating attentional signals from one region to another. Subcortical structures are also heavily involved and influenced by reward and emotion, as discussed in the following section.

One emerging theme is that disparate modes of processing might exist in the different brain regions identified above. Areas such as the LIP and FEF, and subcortical structures such as the SC, might normally operate in a feature-agnostic mode, encoding salience and facilitating or inhibiting regions of visual space according to behavioral goals, but without regard to detailed visual features. (This does not, however, preclude these areas from developing feature selectivity through operant training [16], conditioning [71] or task demands [72]). Conversely, visual cortices (areas V1, V2, V4), the IT and the PFC might be operating in a feature-committed mode, modulating responses depending on the exact visual features that give rise to BU salience and/or TD relevance. The pulvinar might then serve as a bidirectional translator, converting fine-grained, feature-committed TD signals to coarser, feature-agnostic TD signals and vice versa. This dichotomy between feature-agnostic and feature-committed TD signals gives rise to interesting hypotheses about possible mechanisms in which TD attention exerts its influence on neural responses in sensory cortex, and thus affects attentional allocation and gaze behavior.

The role of reward and emotion in TD attention

Until recently, studies of visual attention have traditionally tended to avoid non-visual aspects of cortical and subcortical neuronal responses to manipulations of attention. This has begun to change with a small number of psychophysical and electrophysiological studies that have explored the interplay between reward and attention.

To investigate the role of reward in modulating attention-related responses in the LIP, stimulus selection has been dissociated from motor selection in monkeys [71]. With training, LIP neurons exhibit a strong sustained bias toward the location of a conditioned stimulus, even when a saccade in the opposite direction was required to reveal the outcome of the trial. This suggests that LIP neurons encode 'the value of information' [23] and prioritize spatial locations based on this value.

Studies using operant conditioning paradigms demonstrate effects related to improvements in the volitional TD process. However, learning is also the primary method for augmenting the mandatory TD process. It has been shown that the FEF develops systematic biases, akin to a mandatory TD signal, thereby facilitating shifts of attention in the direction of the feature when it is present at any location [73,74]. More recently, a similar tendency was found in humans performing a visual search task in which the target changed on every trial, which therefore precluded subjects from simply learning a limited set of target features [75]. Subjects' performance improved, demonstrating an improved ability to quickly extract information from a brief preview of the target before each trial, and to then use this information to shape TD signals and guide attention. Learning and reward paradigms can therefore influence ability to both generate TD biasing signals (i.e. volitional TD process) and introduce systematic biases (i.e. mandatory TD process).

Reward plays an important role in modulating attentional signals, and the basal ganglia, which consist of dopaminergic nuclei in the substantia nigra pars reticulata

(SNr), the caudate and the putamen, are essential in encoding reward signals [76]. The basal ganglia are integrally connected to the oculomotor system through the connection of the SC to the SNr [77]. Reward signals (TD) from frontal cortices are transmitted to the caudate, which then inhibits the SNr, which in turn pauses the tonic inhibition from the SNr to the SC, releasing it from inhibition and enabling saccades [78]. This follows a more general scheme in the CNS in which the basal ganglia circuit continually inhibits movement of all limbs until an explicit command to make a motor movement is received from cortical or subcortical regions. Furthermore, it is also possible that reward plays a strong role in influencing a subcortical salience map that can cause instant oculomotor reflexes.

A recent study has shed new light on the SNr to SC connection by demonstrating that SNr fibers connect not only to excitatory neurons in the SC, but also to local GABAergic neurons in the intermediate layers of the SC [79]. Therefore, the SNr is involved in shaping the balance of inhibition and excitation in the local SC circuit. SC involvement in attentional selection and the strong role of the SNr in reward render the SNr-SC connection an important one because in most studies, especially physiological studies in monkeys, paradigms are based on the elements of operant conditioning and reinforcement learning with a crucial role for reward (see [78,80] for more detailed discussions).

Sensory processing is also amenable to modulation by brain regions encoding emotions. In particular, it is known that the amygdala has reciprocal connections with both early and late visual areas and can thus give priority, through modulation, to stimuli of ecological relevance [81]. Using a combination of functional magnetic resonance imaging (fMRI) and a study of lesion patients, it was found that visual areas such as the fusiform gyrus receive input from the amygdala and exhibit enhanced responses to affective stimuli [82]. Such modulation by emotion matches response enhancement observed through attentional allocation. Furthermore, it has been shown that emotional and attentional modulations can act independently, as observed in patients with lesions of the amygdala, whose fusiform cortex exhibited responses modulated by attention but not emotion [82]. Affective stimuli can therefore impinge on sensory signals independently of attention; however, the very enhancement due to emotional valence might render the stimuli salient and thus draw more attention. Attention and emotion might thus act independently on the sensory signals and the behavioral relevance of these sensory inputs might be determined by the cumulative effects of both attention and emotion.

One proposal for neural mechanisms and regions involved in fusion of affective inputs with purely visual aspects driving attention has recently been suggested based on a search task in human subjects using fMRI [83]. The frontoparietal spatial attention network, consisting of the superior parietal lobule (SPL), the inferior parietal lobule (IPL) and the FEF, was activated when the cue was purely spatial. However, when the cue contained both spatial and emotional information, limbic and subcortical structures including the posterior cingulate

cortex (PCC), the amygdala and the orbitofrontal cortex were activated, in addition to the frontoparietal network. This study also found selectivity in the PCC for responding only to cues that had emotional valence [83]. These results suggest that the cingulate gyrus, which receives inputs from the amygdala and sends outputs to the frontoparietal network, might serve as the gateway for affective inputs to fuse with spatial biasing signals. This gives rise to a TD salience map in the frontoparietal network, complete with affective and spatial priority information.

Although evidence remains limited, a number of studies have demonstrated links between the attentional network and reward and emotional centers. Such connections must be taken into account when considering TD networks, because most experimental paradigms involving TD attention to date have used reward and/or emotional valence to train and motivate human or animal participants.

The role of oscillatory activity and neuromodulation in TD attention

It has recently been suggested that synchronous activity (in the gamma range, 50–80 Hz) between cortical regions might serve as the basis for attentional facilitation and cortical computations [84]. In this proposal, neuronal populations representing inputs and decision centers all consist of rhythmically active neural ensembles with distinct excitatory and inhibitory phases. Inhibitory interneurons in each ensemble rhythmically inhibit excitatory pyramidal neurons, thereby establishing a rhythm. Two neural ensembles can then synchronize through phase-locking. This gives rise to a winner-take-all mechanism among two competing inputs feeding into a single higher cortical decision area, through synchronization between the higher area and one selected input. Synchrony between the input and higher areas can be established in a TD or BU manner. In the TD case, a region in higher cortical regions might establish a gamma synchrony with a lower sensory area by phase-locking.

Data from several studies demonstrate that gamma oscillations in the cortex are correlated with attention [45,85]. Disparate brain regions might synchronize their activity in the gamma band when an animal is attending to a particular stimulus. A specific example of this type of coupling is that observed between the FEF and area V4 in monkeys. When attending to a stimulus, coupling through gamma oscillations during attention was observed between neurons in the FEF and V4 [45]. Oscillations in lower frequency bands, such as the alpha and delta bands, have also been implicated in sensory selection [86]. Specifically, in the presence of rhythmic stimuli, delta band oscillations in visual cortex entrain to the rhythm of the stimuli [86]. In doing so, periods of excitability in sensory cortex are aligned with events in the attended stream. In this manner, behaviorally relevant events in the input can be detected more reliably. The same study also showed that the phase of the low-frequency band can modulate amplitudes in higher-frequency bands, such as the gamma band essential for attention. Thus, oscillations in both the gamma and lower-frequency bands are essential neural mechanisms for sensory selection and attention.

The neurochemical basis for attention further supports the notion that synchrony is a possible mechanism for TD

attention. Several studies have described acetylcholine (ACh) as the major neurotransmitter involved in mediating attention at the neuronal level [87]. Using pharmacological manipulations, it was found that attentional modulation in area V1 could be enhanced by low doses of ACh [88]. Furthermore, injection of a muscarinic ACh receptor (mAChR) antagonist eliminated such facilitation, but a nicotinic ACh receptor (nAChR) antagonist did not. This demonstrates that ACh acts through mAChRs to modulate attention. Such modulation might enhance processing in sensory areas, a property of TD attention. It has been demonstrated that pharmacological modulation of glutamatergic transmission in the PFC causes an increase in cholinergic release in the PPC [89]. Given the evidence from the studies discussed above, it is reasonable to hypothesize that one neurochemical process by which the PFC could be involved in TD biasing is modulation of ACh release in sensory areas.

One method that has been suggested for achieving gamma synchrony is the disinhibition of pyramidal cells from inhibitory interneuron activity through cholinergic inputs [90]. This suggests that the cholinergic system might also give rise to the gamma synchrony correlated with attention [84]. Taken together, this evidence suggests that one possible mechanism involved in the selection of relevant sensory stimuli is via modulation of ACh by higher cortical regions, such as the PFC, onto sensory cortical regions, which in turn would induce more powerful gamma synchronies between sensory and higher cortical regions. However, it is currently unclear whether gamma synchrony modulation or firing rate modulation is the core mechanism involved in TD attention. This question was addressed using a biophysically realistic computational model of a single layer of visual cortex receiving attentional inputs [91]. The model of the visual cortex consisted of neurons with glutamatergic synapses. These synapses were modeled with two types of glutamate receptors, AMPA and NMDA. Modulation of the ratio of AMPA to NMDA receptor conductance gave rise to both firing rate and gamma synchrony modulation in an independent manner. This suggests that TD attention might be able to regulate these two systems in an independent manner to set or modify gain in sensory areas. Despite the paucity of conclusive empirical evidence, neural gamma synchrony and the concept of glutamatergic modulation in PFC giving rise to ACh modulation in sensory areas provide a compelling potential neural mechanism for TD attention.

Computational modeling

Physiological studies have guided several theoretical and computational models of attention. Building on the influential feature integration theory [2], guided search theory hypothesizes that massively parallel pre-attentive processes can be guided by TD biasing for features and locations [92]. This theory brings TD elements to a basic BU model of attention [93], which computes individual features at different scales and then combines these features to form a saliency map. A unifying normalization model of attention has recently been proposed and accounts for many effects of TD attention onto visual areas (Figure 6a) [37]. In this model, the neuronal population response of sensory cortex

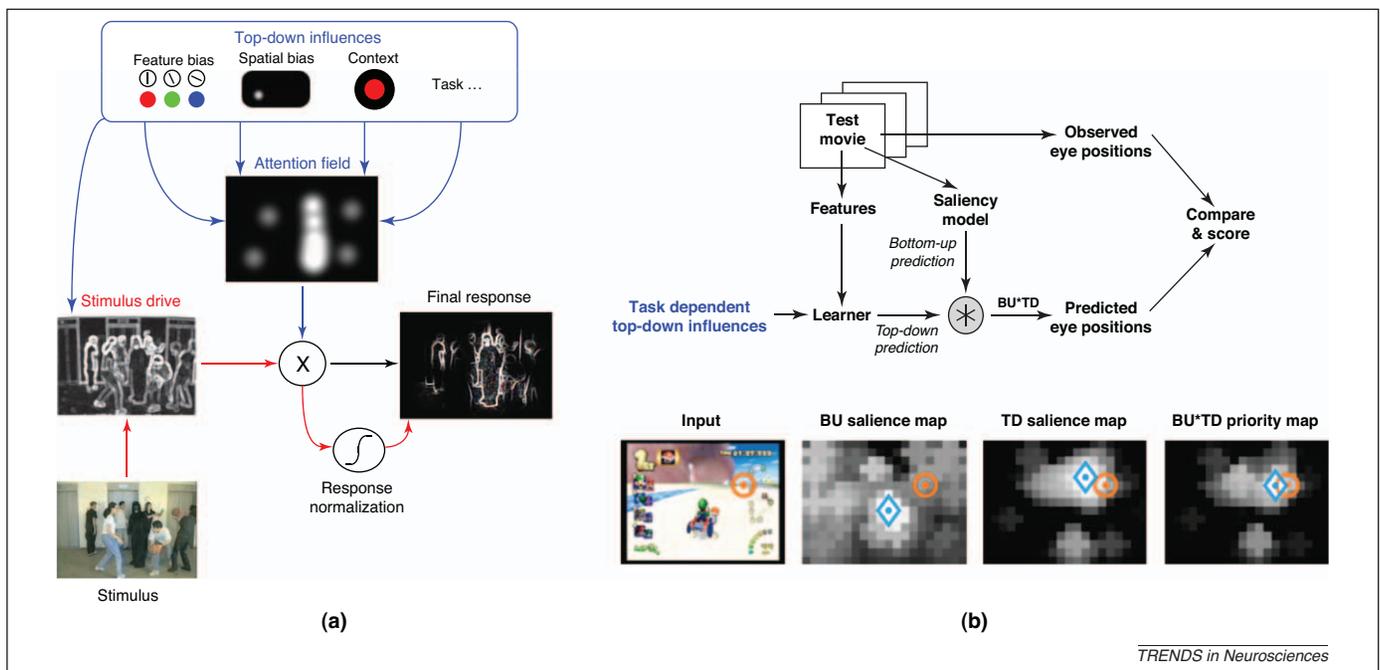


Figure 6. Computational modeling of TD attention. **(a)** Model of attention processing inspired by the normalization model of attention [37]. A visual stimulus can be processed by early visual processing stages and this gives rise to stimulus drive. Stimulus drive can then be combined with an attention field that can provide TD modulation over space. Although the model does not specify how the attention field is formed, we hypothesize that TD influences are responsible for this. Note how some TD signals might directly modulate or shape the stimulus drive (e.g. by shifting receptive fields or affecting orientation preference). After combination with the attentional field, responses undergo divisive normalization and contrast gain control before outputting the final response. Figure adapted with permission from [37]. Example input frame provided by Daniel Simons [138]. **(b)** A model based on related ideas that provides a computer implementation applied to the analysis of human gaze behavior while engaged in complex naturalistic tasks (e.g. driving) [97]. A task-dependent learner component builds, during a training phase, associations between distinct coarse types of scenes and observed eye movements (e.g. drivers tend to look to the right when the road turns right). During testing, exposure to similar scenes gives rise to a TD saliency map (similar to the attention field in (a)), which is further combined with a BU saliency map (similar to the stimulus drive) to give rise to the final BU*TD priority map that guides attention. Blue diamonds represent the peak location for each map and orange circles represent the current eye position of a human player. Figure adapted with permission from [97].

to a stimulus is determined by a competitive normalization process that combines stimulus drive, suppressive drive and an attention field. Although this model successfully captures a wide range of single-unit observations, it does not elucidate how the attention field is formed. This concept is related to the idea of a task-relevance map, a topographical map of visual space that might highlight locations or features of current behavioral relevance and might then act as a mask or filter over the BU saliency map [94]. The task-relevance map might be populated by combining information about desired features (e.g. look for red items), cued spatial locations (e.g. instructions that the target is to the right), scene gist and context (e.g. when looking for a stapler in an office, focus first on desktops), short-term memory of objects and features at previously visited locations, and TD expectations arising from reasoning about what has been discovered so far in light of the task (e.g. if searching for a computer mouse, finding a keyboard and reasoning that the owner of the machine might be right-handed might bias attention to the right of the keyboard) [94]. Interestingly, recent human neuroimaging data provide direct support for such task relevance or TD saliency map possibly located in the intraparietal sulcus (IPS). Indeed, it has been shown that the latter combines, into a single topographic (or, at least, lateralized) map, information about both TD-relevant locations and TD-relevant features [95], and emotional or motivational value of a cued target [83,96]. In a biologically inspired large-scale computer vision implementation, a

similar combination of a TD attention field and BU saliency map was used to predict eye movements of humans engaged in complex tasks (e.g. combat flying or first-person exploration video games) [97]. Given the complexity of these tasks and the multiple interacting TD goals involved, this model did not attempt to fully analyze and recognize all objects in scenes and to assess them in light of the task goals. Instead, the TD map was obtained from learned associations between particular types of scenes (summarized by a simple vector of features capturing their gist) and the locations that humans looked at when engaged in the same task and exposed to similar visual scenes (Figure 6b).

At one extreme, TD attention signals might just consist of a single bit of information – to ‘enhance’ or not – with target visual areas interpreting it in different manners depending on context and on visual inputs. One advantage of such a solution is the low TD communication bandwidth, but an obvious drawback is the inflexibility of signal content. At the other extreme, the brain areas where TD signals originate might address every sensory neuron individually and explicitly modulate the neuron’s activity; for example, increasing gain by some specific amount, sharpening tuning, and increasing baseline activity. Such a scheme would afford maximal flexibility, but at the cost of both enormous TD communication bandwidth and high computational requirements in areas where TD signals originate, to compute the exact values for all these signals. The true nature of TD signals is likely to lie between these two extremes, as further elaborated below.

The Guided Search model lies towards the low-bandwidth end of the spectrum, with TD signals imposing spatial attention modulation over coarse regions of visual space and coarse visual features (e.g. a single TD attention weight for each of red, green, blue or yellow colors, or steep, shallow, left or right orientations) [92]. Two recent studies have refined this proposal. First, in human eye-tracking experiments it has been shown that attention and gaze can effectively be guided towards rather fine sub-bands of basic visual features, such as mid-luminance items among low- and high-luminance items, and similarly for size and color saturation [98]. Furthermore, these results have been formalized with a signal-to-noise ratio (SNR)-maximizing model for feature search, whereby the TD gain applied to each sensory neuron is proportional to its ability to distinguish the target of behavioral interest from background clutter [99]. Taken together, these two studies suggest that the bandwidth or granularity of TD signals is unlikely to be extremely low, but rather might consist of at least a few bits for each fine-grained feature sub-band, sufficient to convey optimal biases from the top down. The bandwidth (and number of descending connections) might be higher if different biases can be communicated to different locations of sensory space. At the high extreme, the aforementioned normalization model of attention assumes a highly detailed attention field over space and features [37], implying high-bandwidth TD signals.

Beyond the nature and bandwidth of information conveyed from the top down, computational models have proposed a number of connectivity styles that might be embodied in the biological reality of TD connections. On the one hand, one model has identified a specific dedicated structure (the pulvinar) as a hub or relay for TD signals to reach target visual areas [69]. On the other hand, a more distributed model suggests that TD signals are embedded within the visual areas themselves [100]. In this model, a stimulus is selected at the top level based on an initial sweep of feed-forward information. The spatial selection signals then propagate back and tune lower levels of the (cortical) visual processing hierarchy through a cascade of winner-take-all mechanisms. This view involves retrograde propagation of signals over the processing hierarchy as opposed to direct connections (or through one or a few relays) between top and bottom. A number of models also give specific roles to direct or indirect connections among different levels of the hierarchy, for example between the PFC, FEF, TE and V4 [101]. These models are important because they develop hypotheses for the meaning of large-scale connectivity between brain areas, and these are beginning to be explicitly tested in biological networks using graph-theoretic analyses [102]. Nevertheless, there is a clear lack of specific computational (and experimental) studies that systematically investigate the granularity, bandwidth and specific wiring of TD signals.

Finally, computational theories and models have started to provide hypotheses for the meaning of TD signals. For example, models based on feedback connections from higher cortical areas have been placed in a Bayesian framework, with the suggestion of a generative model that produces a hypothesis about a percept (the prior), then combines this with evidence from BU information to make

Box 1. Outstanding questions

- What is the bandwidth of the TD signal transmitted from one region of the brain to the next? Figure 1 illustrates the two types of signal. A narrow-bandwidth signal (yellow arrow) defines single weights for individual features, whereas a broadband signal (blue arrow) defines the distribution of gain and tuning over the feature space, as well as the interactions within a feature dimension.
- Are TD signals relayed to visual areas through a central hub (e.g. the pulvinar) or does a more distributed mechanism reflect the reality of communication of TD signals to sensory areas?
- What is the representation or encoding of TD signals? In concrete terms, how are behavioral goals represented and communicated to sensory neurons that are tuned to specific features?
- What (if any) computations take place subcortically, independent of the cortex, that would influence attention modulation of sensory perception?

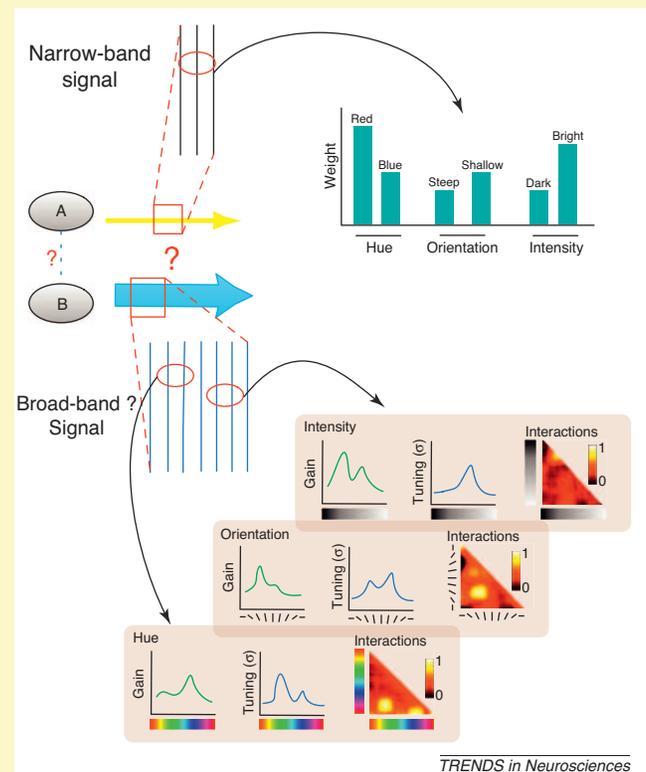


Figure 1. Narrow-band versus broad-band TD biasing signals. The TD biasing signals transmitted from one area 'A' of the brain to another area 'B' can be either narrow-band (yellow arrow) or broad-band (blue arrow) in nature. Narrow-band signals consist of a small set of weights that bias feature preferences in a coarse manner. The bar graph shows a signal that applies a higher gain to neurons tuned to red rather than blue in the color feature dimension, neurons tuned to shallow rather than steep orientations, and neurons tuned to brighter rather than darker stimuli. Broad-band biasing signals (bottom) contain a greater amount of information and might facilitate biasing of features in a detailed manner, weighing gain, tuning and feature interactions independently. Rather than simply setting a weight along a feature dimension, as is the case in the narrow-band example, broad-band TD signals might set a biasing profile along the feature dimension, as shown in the example graphs. Green curves show a biasing profile for gains of neurons along a feature dimension; a peak here would indicate a bias or shift of tuning of neurons for the particular feature value. The interaction triangles on the right show biases for feature interactions. For example, along the hue dimension, there are two hot spots, one indicating a preference for simultaneous occurrence of yellow and red hues and another indicating a preference for red and blue hues.

a final decision on the percept [103]. This approach has been formalized in a hierarchical Bayesian framework [104]. Although these ideas have so far been explored more in the context of the mandatory TD process, they can also

be placed in the context of the volitional TD process. Volitional TD control could then be understood as updating, biasing or disambiguating the prior based on high-level tasks, contextual cues or behavioral goals. In computer-vision models using these principles, it has indeed been shown that TD attention provides great benefit over pure BU processing [105]. For example, TD information can more effectively guide visual search for specific objects in natural scenes (e.g. pedestrians in street scenes) by limiting the search to spatial locations of high prior or posterior probabilities [106–108]. Although computational models have made some headway in both incorporating experimental data and generating predictions to guide further experiments, much remains to be done both experimentally and theoretically to unravel the mechanisms by which TD attentional mechanisms influence BU processing (Box 1).

Conclusion

Attention modulates sensory signals early in the process, exerting its influence on the SC and the thalamus before further modulating signals in cortex. The cumulative effects of this modulation based on both TD and BU influences might be represented by a priority map over visual space. Although there is some debate about the exact locus of the priority map, it is clear that the LIP, FEF and SC exhibit properties that are compatible with the existence of a spatial map encoding behavioral relevance of spatial locations. These three regions might jointly compute or host such a map that is agnostic to the features that caused the priority. Thus, the map fuses both BU and TD influences and drives motor output.

Higher cortical areas such as the PFC send detailed TD signals to sensory areas for biasing of spatial and non-spatial features. Such signals fuse together with reward-related and emotional signals to form the TD influence on attention, which might be reflected in the priority map. Subcortical regions, through their close connection to the reward systems in the brain and their coupling with motor systems, exert strong influences on attentional signals, in addition to being major targets of attentional modulation for motor output. Feedback connections are both pervasive and crucial for the transmission of biasing signals emanating from higher brain regions, especially the frontal cortices that are involved in working memory processes and send descending reward signals. Computational studies highlight the important constraints on the nature, granularity, bandwidth and connectivity style of TD connections. There is a pressing need to build models that take into account physiological data, particularly from microstimulation and lesion studies, which could help to determine the contributions of specific areas to the computations necessary for attentional guidance.

Although the exact mechanisms of TD attention have yet to be completely delineated, there are sufficient data available to demonstrate that attention is mediated by the merging of TD and BU information. As William James eloquently stated, ‘The attentive process, therefore, at its maximum may be physiologically symbolized, by a brain-cell played on in two ways from without and from within’ [109].

Acknowledgements

This work was supported by the Defense Advanced Research Projects Agency (government contract no. HR0011-10-C-0034), the National Science Foundation (CRCNS grant number BCS-0827764), General Motors Corporation, and the Army Research Office (grant no. W911NF-08-1-0360). The authors affirm that the views expressed herein are solely their own, and do not represent the views of the United States government or any agency thereof. We would also like to thank Robert Desimone, Jack Gallant, Jacqueline Gottlieb and the anonymous reviewers for their helpful comments and suggestions.

References

- James, W. (1890) *The Principles of Psychology*, Harvard University Press
- Treisman, A. and Gelade, G. (1980) A feature-integration theory of attention. *Cogn. Psychol.* 12, 97–136
- Wolfe, J.M. and Horowitz, T.S. (2004) What attributes guide the deployment of visual attention and how do they do it? *Nat. Rev. Neurosci.* 5, 495–501
- Itti, L. and Koch, C. (2001) Computational modelling of visual attention. *Nat. Rev. Neurosci.* 2, 194–203
- Gilbert, C. and Sigman, M. (2007) Brain states: top-down influences in sensory processing. *Neuron* 54, 677–696
- Chun, M.M. and Jiang, Y. (1998) Contextual cueing: implicit learning and memory of visual context guides spatial attention. *Cogn. Psychol.* 36, 28–71
- Desimone, R. and Duncan, J. (1995) Neural mechanisms of selective visual attention. *Annu. Rev. Neurosci.* 18, 193–222
- Noudoost, B. *et al.* (2010) Top-down control of visual attention. *Curr. Opin. Neurobiol.* 20, 183–190
- Miller, E.K. and Cohen, J.D. (2001) An integrative theory of prefrontal cortex function. *Annu. Rev. Neurosci.* 24, 167–202
- Corbetta, M. and Shulman, G.L. (2002) Control of goal-directed and stimulus-driven attention in the brain. *Nat. Rev. Neurosci.* 3, 201–215
- Werblin, F. *et al.* (2001) Parallel processing in the mammalian retina: lateral and vertical interactions across stacked representations. *Prog. Brain Res.* 131, 229–238
- Goodale, M.A. and Milner, A.D. (1992) Separate visual pathways for perception and action. *Trends Neurosci.* 15, 20–25
- Wurtz, R. and Mohler, C. (1976) Enhancement of visual responses in monkey striate cortex and frontal eye fields. *J. Neurophysiol.* 39, 766–772
- Fecteau, J. and Munoz, D. (2006) Saliency, relevance, and firing: a priority map for target selection. *Trends Cogn. Sci.* 10, 382–390
- Bisley, J.W. and Goldberg, M.E. (2010) Attention, intention, and priority in the parietal lobe. *Annu. Rev. Neurosci.* 33, 1–21
- Thompson, K.G. and Bichot, N.P. (2005) A visual saliency map in the primate frontal eye field. *Prog. Brain Res.* 147, 251–262
- Shipp, S. (2003) The functional logic of cortico-pulvinar connections. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 358, 1605–1624
- McAlonan, K. *et al.* (2008) Guarding the gateway to cortex with attention in visual thalamus. *Nature* 456, 391–394
- O’Connor, D. *et al.* (2002) Attention modulates responses in the human lateral geniculate nucleus. *Nat. Neurosci.* 5, 1203–1209
- Buffalo, E.A. *et al.* (2010) A backward progression of attentional effects in the ventral stream. *Proc. Natl. Acad. Sci. U.S.A.* 107, 361–365
- McAdams, C.J. and Maunsell, J.H. (1999) Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *J. Neurosci.* 19, 431–441
- Martinez-Trujillo, J.C. and Treue, S. (2004) Feature-based attention increases the selectivity of population responses in primate visual cortex. *Curr. Biol.* 14, 744–751
- Gottlieb, J. and Balan, P. (2010) Attention as a decision in information space. *Trends Cogn. Sci.* 14, 240–248
- Gottlieb, J. (2007) From thought to action: the parietal cortex as a bridge between perception, action, and cognition. *Neuron* 53, 9–16
- Gottlieb, J. *et al.* (1998) The representation of visual saliency in monkey parietal cortex. *Nature* 391, 481–484
- Bisley, J. and Goldberg, M. (2003) Neuronal activity in the lateral intraparietal area and spatial attention. *Science* 299, 81–86
- Ungerleider, L.G. *et al.* (2008) Cortical connections of area V4 in the macaque. *Cereb. Cortex* 18, 477–499

- 28 Lamme, V.A. and Roelfsema, P.R. (2000) The distinct modes of vision offered by feedforward and recurrent processing. *Trends Neurosci.* 23, 571–579
- 29 Monosov, I. *et al.* (2008) Measurements of simultaneously recorded spiking activity and local field potentials suggest that spatial selection emerges in the frontal eye field. *Neuron* 57, 614–625
- 30 Buschman, T. and Miller, E. (2007) Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. *Science* 315, 1860–1862
- 31 Lebedev, M.A. *et al.* (2004) Representation of attended versus remembered locations in prefrontal cortex. *PLoS Biol.* 2, e365
- 32 Luck, S.J. *et al.* (1997) Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *J. Neurophysiol.* 77, 24–42
- 33 Bichot, N.P. *et al.* (2005) Parallel and serial neural mechanisms for visual search in macaque area V4. *Science* 308, 529–534
- 34 Saenz, M. *et al.* (2003) Global feature-based attention for motion and color. *Vis. Res.* 43, 629–637
- 35 Treue, S. and Martinez Trujillo, J.C. (1999) Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* 399, 575–579
- 36 Motter, B.C. (1994) Neural correlates of attentive selection for color or luminance in extrastriate area V4. *J. Neurosci.* 14, 2178–2189
- 37 Reynolds, J.H. and Heeger, D.J. (2009) The normalization model of attention. *Neuron* 61, 168–185
- 38 Reynolds, J. and Chelazzi, L. (2004) Attentional modulation of visual processing. *Annu. Rev. Neurosci.* 27, 611–647
- 39 Knudsen, E.I. (2007) Fundamental components of attention. *Annu. Rev. Neurosci.* 30, 57–78
- 40 Tomita, H. *et al.* (1999) Top-down signal from prefrontal cortex in executive control of memory retrieval. *Nature* 401, 699–703
- 41 Rossi, A. *et al.* (2007) Top down attentional deficits in macaques with lesions of lateral prefrontal cortex. *J. Neurosci.* 27, 11306–11314
- 42 Opris, I. *et al.* (2005) Microstimulation of the dorsolateral prefrontal cortex biases saccade target selection. *J. Cogn. Neurosci.* 17, 893–904
- 43 Moore, T. and Armstrong, K. (2003) Selective gating of visual signals by microstimulation of frontal cortex. *Nature* 421, 370–373
- 44 Winkowski, D. and Knudsen, E. (2008) Distinct mechanisms for top-down control of neural gain and sensitivity in the owl optic tectum. *Neuron* 60, 698–708
- 45 Gregoriou, G.G. *et al.* (2009) High-frequency, long-range coupling between prefrontal and visual cortex during attention. *Science* 324, 1207–1210
- 46 Wardak, C. *et al.* (2006) Contribution of the monkey frontal eye field to covert visual attention. *J. Neurosci.* 26, 4228–4235
- 47 Cutrell, E. and Marrocco, R. (2002) Electrical microstimulation of primate posterior parietal cortex initiates orienting and alerting components of covert attention. *Exp. Brain Res.* 144, 103–113
- 48 Balan, P.F. and Gottlieb, J. (2009) Functional significance of nonspatial information in monkey lateral intraparietal area. *J. Neurosci.* 29, 8166–8176
- 49 Friedman-Hill, S. *et al.* (2003) Posterior parietal cortex and the filtering of distractors. *Proc. Natl. Acad. Sci. U.S.A.* 100, 4263–4268
- 50 Saalmann, Y. *et al.* (2007) Neural mechanisms of visual attention: how top-down feedback highlights relevant locations. *Science* 316, 1612
- 51 Liu, Y. *et al.* (2010) Intention and attention: different functional roles for LIPD and LIPV. *Nat. Neurosci.* 13, 495–500
- 52 Sherman, S. (2007) The thalamus is more than just a relay. *Curr. Opin. Neurobiol.* 17, 417–422
- 53 Kawasaki, K. and Sheinberg, D.L. (2008) Learning to recognize visual objects with microstimulation in inferior temporal cortex. *J. Neurophysiol.* 100, 197–211
- 54 Afraz, S-R. *et al.* (2006) Microstimulation of inferotemporal cortex influences face categorization. *Nature* 442, 692–695
- 55 Cavanaugh, J. and Wurtz, R. (2004) Subcortical modulation of attention counters change blindness. *J. Neurosci.* 24, 11236–11243
- 56 Muller, J. *et al.* (2005) Microstimulation of the superior colliculus focuses attention without moving the eyes. *Proc. Natl. Acad. Sci. U.S.A.* 102, 524–529
- 57 Carello, C. and Krauzlis, R., (2004) Manipulating intent evidence for a causal role of the superior colliculus in target selection. *Neuron* 43, 575–583
- 58 McPeck, R. and Keller, E. (2004) Deficits in saccade target selection after inactivation of superior colliculus. *Nat. Neurosci.* 7, 757–763
- 59 Ignashchenkova, A. *et al.* (2003) Neuron-specific contribution of the superior colliculus to overt and covert shifts of attention. *Nat. Neurosci.* 7, 56–64
- 60 Lovejoy, L. and Krauzlis, R. (2009) Inactivation of primate superior colliculus impairs covert selection of signals for perceptual judgments. *Nat. Neurosci.* 13, 261–266
- 61 Kaas, J.H. and Lyon, D.C. (2007) Pulvinar contributions to the dorsal and ventral streams of visual processing in primates. *Brain Res. Rev.* 55, 285–296
- 62 Leh, S. *et al.* (2008) The connectivity of the human pulvinar: a diffusion tensor imaging tractography study. *Int. J. Biomed. Imaging* 2008, 1–5
- 63 Robinson, D.L. and Petersen, S.E. (1992) The pulvinar and visual salience. *Trends Neurosci.* 15, 127–132
- 64 Desimone, R. *et al.* (1990) Attentional control of visual perception: cortical and subcortical mechanisms. *Cold Spring Harb. Symp. Quant. Biol.* 55, 963–971
- 65 Snow, J. *et al.* (2009) Impaired attentional selection following lesions to human pulvinar: evidence for homology between human and monkey. *Proc. Natl. Acad. Sci. U.S.A.* 106, 4054–4059
- 66 Arend, I. *et al.* (2008) Spatial and temporal deficits are regionally dissociable in patients with pulvinar lesions. *Brain* 131, 2140–2152
- 67 Bender, D.B. and Baizer, J.S. (1990) Saccadic eye movements following kainic acid lesions of the pulvinar in monkeys. *Exp. Brain Res.* 79, 467–478
- 68 Bender, D.B. and Butter, C.M. (1987) Comparison of the effects of superior colliculus and pulvinar lesions on visual search and tachistoscopic pattern discrimination in monkeys. *Exp. Brain Res.* 69, 140–154
- 69 Olshausen, B.A. *et al.* (1993) A neurobiological model of visual attention and invariant pattern recognition based on dynamic routing of information. *J. Neurosci.* 13, 4700–4719
- 70 Berman, R.A. and Wurtz, R.H. (2010) Functional identification of a pulvinar path from superior colliculus to cortical area MT. *J. Neurosci.* 30, 6342–6354
- 71 Peck, C.J. *et al.* (2009) Reward modulates attention independently of action value in posterior parietal cortex. *J. Neurosci.* 29, 11182–11191
- 72 Toth, L.J. and Assad, J.A. (2002) Dynamic coding of behaviourally relevant stimuli in parietal cortex. *Nature* 415, 165–168
- 73 Bichot, N.P. and Schall, J.D. (1999) Effects of similarity and history on neural mechanisms of visual selection. *Nat. Neurosci.* 2, 549–554
- 74 Bichot, N.P. *et al.* (1996) Visual feature selectivity in frontal eye fields induced by experience in mature macaques. *Nature* 381, 697–699
- 75 Baluch, F. and Itti, L. (2010) Training top-down attention improves performance on a triple-conjunction search task. *PLoS ONE* 5, e9127
- 76 Graybiel, A. (2005) The basal ganglia: learning new tricks and loving it. *Curr. Opin. Neurobiol.* 15, 638–644
- 77 Boehnke, S.E. and Munoz, D.P. (2008) On the importance of the transient visual response in the superior colliculus. *Curr. Opin. Neurobiol.* 18, 544–551
- 78 Hikosaka, O. *et al.* (2006) Basal ganglia orient eyes to reward. *J. Neurophysiol.* 95, 567–584
- 79 Kaneda, K. *et al.* (2008) Nigral inhibition of GABAergic neurons in mouse superior colliculus. *J. Neurosci.* 28, 11071–11078
- 80 Shires, J. *et al.* (2010) Shedding new light on the role of the basal ganglia-superior colliculus pathway in eye movements. *Curr. Opin. Neurobiol.* 20, 717–725
- 81 Pessoa, L. (2008) On the relationship between emotion and cognition. *Nat. Rev. Neurosci.* 9, 148–158
- 82 Vuilleumier, P. *et al.* (2004) Distant influences of amygdala lesion on visual cortical activation during emotional face processing. *Nat. Neurosci.* 7, 1271–1278
- 83 Mohanty, A. *et al.* (2009) Search for a threatening target triggers limbic guidance of spatial attention. *J. Neurosci.* 29, 10563–10572
- 84 Fries, P. (2009) Neuronal gamma-band synchronization as a fundamental process in cortical computation. *Annu. Rev. Neurosci.* 32, 209–224
- 85 Fries, P. *et al.* (2002) Oscillatory neuronal synchronization in primary visual cortex as a correlate of stimulus selection. *J. Neurosci.* 22, 3739–3754

- 86 Lakatos, P. *et al.* (2008) Entrainment of neuronal oscillations as a mechanism of attentional selection. *Science* 320, 110–113
- 87 Sarter, M. *et al.* (2005) Unraveling the attentional functions of cortical cholinergic inputs: interactions between signal-driven and cognitive modulation of signal detection. *Brain Res. Rev.* 48, 98–111
- 88 Herrero, J.L. *et al.* (2008) Acetylcholine contributes through muscarinic receptors to attentional modulation in V1. *Nature* 454, 1110–1114
- 89 Nelson, C.L. *et al.* (2005) Prefrontal cortical modulation of acetylcholine release in posterior parietal cortex. *Neuroscience* 132, 347–359
- 90 Deco, G. and Thiele, A. (2009) Attention: oscillations and neuropharmacology. *Eur. J. Neurosci.* 30, 347–354
- 91 Buehlmann, A. and Deco, G. (2008) The neuronal basis of attention: rate versus synchronization modulation. *J. Neurosci.* 28, 7679–7686
- 92 Wolfe, J. (1994) Guided search 2.0. A revised model of visual search. *Psychonom. Bull. Rev.* 1, 202–238
- 93 Koch, C. and Ullman, S. (1985) Shifts in selective visual attention: towards the underlying neural circuitry. *Hum. Neurobiol.* 4, 219–227
- 94 Navalpakkam, V. and Itti, L. (2005) Modeling the influence of task on attention. *Vis. Res.* 45, 205–231
- 95 Egner, T. *et al.* (2008) Neural integration of top-down spatial and feature-based information in visual search. *J. Neurosci.* 28, 6141–6151
- 96 Mohanty, A. *et al.* (2008) The spatial attention network interacts with limbic and monoaminergic systems to modulate motivation-induced attention shifts. *Cereb. Cortex* 18, 2604–2613
- 97 Peters, R. and Itti, L. (2007) Beyond bottom-up: incorporating task-dependent influences into a computational model of spatial attention. In *Proceedings of the 2007 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1–8
- 98 Navalpakkam, V. and Itti, L. (2006) Top-down attention selection is fine-grained. *J. Vis.* 6, 1180–1193
- 99 Navalpakkam, V. and Itti, L. (2007) Search goal tunes visual features optimally. *Neuron* 53, 605–617
- 100 Tsotsos, J.K. *et al.* (1995) Modeling visual-attention via selective tuning. *Artif. Intell.* 78, 507–545
- 101 Hamker, F.H. (2006) Modeling feature-based attention as an active top-down inference process. *Biosystems* 86, 91–99
- 102 Bullmore, E. and Sporns, O. (2009) Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat. Rev. Neurosci.* 10, 186–198
- 103 Gregory, R. (1970) *The Intelligent Eye*, McGraw-Hill
- 104 Lee, T. and Mumford, D. (2003) Hierarchical Bayesian inference in the visual cortex. *J. Opt. Soc. Am. A* 20, 1434–1448
- 105 Frintrop, S. *et al.* (2010) Computational visual attention systems and their cognitive foundation: a survey. *ACM Trans. Appl. Percept.* 7, 6
- 106 Najemnik, J. and Geisler, W.S. (2005) Optimal eye movement strategies in visual search. *Nature* 434, 387–391
- 107 Torralba, A. *et al.* (2006) Contextual guidance of eye movements and attention in real-world scenes: the role of global features in object search. *Psychol. Rev.* 113, 766–786
- 108 Ehinger, K.A. *et al.* (2009) Modeling search for people in 900 scenes: a combined source model of eye guidance. *Vis. Cogn.* 17, 945–978
- 109 James, W. (1892) *Talks to Teachers on Psychology and to Students on Some of Life's Ideals*, Adamant Media Corporation
- 110 Rafal, R.D. and Posner, M.I. (1987) Deficits in human visual spatial attention following thalamic lesions. *Proc. Natl. Acad. Sci. U.S.A.* 84, 7349–7353
- 111 Pezaris, J.S. and Reid, R.C. (2007) Demonstration of artificial visual percepts generated through thalamic microstimulation. *Proc. Natl. Acad. Sci. U.S.A.* 104, 7670–7675
- 112 Schmid, M.C. *et al.* (2010) Blindsight depends on the lateral geniculate nucleus. *Nature* 466, 373–377
- 113 Page, W.K. *et al.* (1994) Magnocellular or parvocellular lesions in the lateral geniculate nucleus of monkeys cause minor deficits of smooth pursuit eye movements. *Vis. Res.* 34, 223–239
- 114 Tehovnik, E.J. *et al.* (2003) Differential effects of laminar stimulation of V1 cortex on target selection by macaque monkeys. *Eur. J. Neurosci.* 16, 751–760
- 115 Tehovnik, E.J. *et al.* (2003) Saccadic eye movements evoked by microstimulation of striate cortex. *Eur. J. Neurosci.* 17, 870–878
- 116 Murphey, D.K. and Maunsell, J.H.R. (2007) Behavioral detection of electrical microstimulation in different cortical visual areas. *Curr. Biol.* 17, 862–867
- 117 Moore, T. *et al.* (2001) Direction of motion discrimination after early lesions of striate cortex (V1) of the macaque monkey. *Proc. Natl. Acad. Sci. U.S.A.* 98, 325–330
- 118 Yoshida, M. *et al.* (2008) Striate cortical lesions affect deliberate decision and control of saccade: implication for blindsight. *J. Neurosci.* 28, 10517–10530
- 119 Buffalo, E. *et al.* (2005) Impaired filtering of distracter stimuli by TE neurons following V4 and TEO lesions in macaques. *Cereb. Cortex* 15, 141–151
- 120 Bachevalier, J. and Mishkin, M. (1994) Effects of selective neonatal temporal lobe lesions on visual recognition memory in rhesus monkeys. *J. Neurosci.* 14, 2128–2139
- 121 Petrides, M. (2000) Dissociable roles of mid-dorsolateral prefrontal and anterior inferotemporal cortex in visual working memory. *J. Neurosci.* 20, 7496–7503
- 122 Nichols, M.J. and Newsome, W.T. (2002) Middle temporal visual area microstimulation influences veridical judgments of motion direction. *J. Neurosci.* 22, 9530–9540
- 123 Bisley, J.W. *et al.* (2001) Microstimulation of cortical area MT affects performance on a visual working memory task. *J. Neurophysiol.* 85, 187–196
- 124 Newsome, W.T. and Paré, E.B. (1988) A selective impairment of motion perception following lesions of the middle temporal visual area (MT). *J. Neurosci.* 8, 2201–2211
- 125 Hanks, T.D. *et al.* (2006) Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. *Nat. Neurosci.* 9, 682–689
- 126 Moore, T. and Fallah, M. (2004) Microstimulation of the frontal eye field and its effects on covert spatial attention. *J. Neurophysiol.* 91, 152–162
- 127 Juan, C. *et al.* (2004) Dissociation of spatial attention and saccade preparation. *Proc. Natl. Acad. Sci. U.S.A.* 101, 15541–15544
- 128 Ruff, C.C. *et al.* (2006) Concurrent TMS-fMRI and psychophysics reveal frontal influences on human retinotopic visual cortex. *Curr. Biol.* 16, 1479–1488
- 129 Neggers, S.F.W. *et al.* (2007) TMS pulses on the frontal eye fields break coupling between visuospatial attention and eye movements. *J. Neurophysiol.* 98, 2765–2778
- 130 Wegener, S.P. *et al.* (2008) Microstimulation of monkey dorsolateral prefrontal cortex impairs antisaccade performance. *Exp. Brain Res.* 190, 463–473
- 131 Rasmusson, D.D. *et al.* (2007) Inactivation of prefrontal cortex abolishes cortical acetylcholine release evoked by sensory or sensory pathway stimulation in the rat. *Neuroscience* 149, 232–241
- 132 Rudolph, K. and Pasternak, T. (1999) Transient and permanent deficits in motion perception after lesions of cortical areas MT and MST in the macaque monkey. *Cereb. Cortex* 9, 90–100
- 133 De Weerd, P. *et al.* (1999) Loss of attentional stimulus selection after extrastriate cortical lesions in macaques. *Nat. Neurosci.* 2, 753–758
- 134 Gregory, R. (1997) Knowledge in perception and illusion. *Philos. Trans. R. Soc. B Biol. Sci.* 352, 1121–1127
- 135 Van Essen, D.C. (2002) Windows on the brain: the emerging role of atlases and databases in neuroscience. *Curr. Opin. Neurobiol.* 12, 574–579
- 136 Van Essen, D.C. *et al.* (2001) An integrated software suite for surface-based analyses of cerebral cortex. *J. Am. Med. Inform. Assoc.* 8, 443–459
- 137 Pascual-Leone, A. and Walsh, V. (2001) Fast backprojections from the motion to the primary visual area necessary for visual awareness. *Science* 292, 510–512
- 138 Simons, D.J. and Chabris, C.F. (1999) Gorillas in our midst: sustained inattentive blindness for dynamic events. *Perception* 28, 1059–1074