



Published in final edited form as:

Neurobiol Learn Mem. 2009 February ; 91(2): 121–128. doi:10.1016/j.nlm.2008.09.013.

Parietal control of attentional guidance: the significance of sensory, motivational and motor factors

Jacqueline Gottlieb^{1,2}, Puiu Balan¹, Jeff Oristaglio¹, and Mototaka Suzuki¹

¹ Department of Neuroscience, Columbia University

² Department of Psychiatry, Columbia University

Abstract

The lateral intraparietal area (LIP), a portion of monkey posterior parietal cortex, has been implicated in spatial attention. We review recent evidence showing that LIP encodes a priority map of the external environment that specifies the momentary locus of attention and is activated in a variety of behavioral tasks. The priority map in LIP is shaped by task-specific motor, cognitive and motivational variables, the functional significance of which is not entirely understood. We suggest that these modulations represent teaching signals by which the brain learns to identify attentional priority to stimuli based on the task-specific associations between these stimuli, the required decision and expected outcome.

Introduction

Judicious selection is at the heart of goal directed behavior. To select appropriately in a complex environment an organism must solve two problems. First, the organism must identify the most relevant source of information from among many potential alternatives. Second, the organism must be able to focus on the selected object despite interference from irrelevant distractions. These two operations – identifying and focusing on specific items – are closely related to each other and are solved by a wide family of mechanisms collectively known as selective attention. It is often claimed that attention is required for overcoming capacity limitations inherent in neural processing: because sensory systems receive more information than they can process, attention provides the mean of prioritizing and limiting the amount of information that must be processed at any given time. However, the need for appropriate selection remains even in simple environments that do not seriously tax capacity limitations. The essence of attention is therefore the act of assigning credit, or identifying the sources of information that are most relevant in a given context. Generally this decision requires *learning* about the statistical contingencies between various objects, and outcomes. Attention, therefore, must be a dynamic selection mechanism that is exquisitely sensitive to immediate task demands.

Neurophysiological investigations in non-human primates have focused on three areas as being important for the control of attention: the frontal eye field in the frontal lobe, the superior colliculus in the midbrain, and the lateral intraparietal area (LIP) in the parietal cortex. In this

Address for correspondence: Jacqueline Gottlieb, PhD, Department of Neuroscience, Columbia University, 1051 Riverside Drive, Kolb Research Annex, Rm. 569, New York, NY 10032, Phone: 212-543-6931, ext. 500, Fax: 212-543-5816, Email: E-mail: jg2141@columbia.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

review we describe the state of our knowledge about one node of this network, area LIP, which has been especially well investigated and provides an excellent model system for further inquiry into the mechanisms of attention. Converging evidence shows that LIP provides a topographic priority representation of the world that guides the moment by moment locus of attention based on a wide variety of stimulus-related (bottom-up) and task-related (top-down) factors. Most intriguingly, recent evidence shows that this priority map is not stereotyped but takes on task-specific properties – i.e., it appears to be plastic and adaptable to task demands. The significance of these modulations is by and large not well understood. However, we suggest that these modulations may represent teaching signals through which the brain learns to assign attentional priority based on the significance of stimuli for a specific action or a specific outcome. We end by describing a computational model (Roelfsema and van Ooyen, 2005) that appears especially useful for formalizing inquiry into the links between learning and attention.

Methods

General methods and behavioral tasks—Data were collected with standard behavioral and neurophysiological techniques as described previously (Balan and Gottlieb, 2006; Oristaglio et al., 2006). All methods were approved by the Animal Care and Use Committees of Columbia University and New York State Psychiatric Institute as complying with the guidelines within the Public Health Service Guide for the Care and Use of Laboratory Animals. During experimental sessions monkeys sat in a primate chair with their heads were fixed in the straight ahead position. Visual stimuli were presented on a SONY GDM-FW9000 Trinitron monitor (30.8 by 48.2 cm viewing area) located 57 cm in front of the monkeys' eyes.

Identification of LIP—Structural MRI was used to verify that electrode tracks coursed through the lateral bank of the intraparietal sulcus. Before testing on the search task each neuron was first characterized with the memory-saccade task on which, after the monkey fixated a central fixation point, a small annulus (1 deg diameter) was flashed for 100 ms at a peripheral location and, after a brief delay the monkey was rewarded for making a saccade to the remembered location of the annulus. All the neurons described here had significant spatial selectivity in the memory saccade task (1-way Kruskal-Wallis analysis of variance, $p < 0.05$) and virtually all (97%) showed this selectivity during the delay or presaccadic epochs (400–900 ms after target onset and 200 ms before saccade onset).

Covert search task—The basic variant of the covert search task (Figure 2a) was tested with display size of 4 elements. Individual stimuli were scaled with retinal eccentricity and ranged from 1.5 to 3.0 deg in height and 1.0 to 2.0 deg in width. To begin a trial, monkeys fixated a central fixation spot (presented anew on each trial) and grabbed two response bars (Figure 2). Two line elements were then removed from each placeholder, yielding a display with one cue (a right- or left-facing letter “E”) and several unique distractors. Monkeys were rewarded for reporting cue orientation by releasing the right bar for a right-facing cue or the left bar for a left-facing cue within 100–1000 ms of the display change. A correct response was rewarded with a drop of juice, after which the fixation point was removed and the placeholder display was restored. Fixation was continuously enforced to within 1.5–2 deg of the fixation point. Errors (fixation breaks, incorrect, early or late bar releases) were aborted without reward.

In the perturbation version of the search task the initial fixation interval was lengthened to 800 or 1200 ms on each trial and a 50 ms perturbation was presented starting 200 ms before presentation of the target display. To increase task difficulty display size was increased to 12 elements and only a fraction of each line segment was removed from each placeholder. The location of perturbation and target were randomly selected from among a restricted neighborhood of 2 or 3 elements centered in and opposite the neuron's receptive field (RF),

with a spatial relationship determined by behavioral context. Contexts were run in interleaved blocks of ~300 trials. Within each context the location of target and perturbation and 2–3 of the possible perturbation types (see text) were randomly interleaved.

Data analysis—Firing rates were measured from the raw spike times and, unless otherwise stated, statistical tests are based on the Wilcoxon rank test or paired-rank test, or on non-parametric analysis of variance, evaluated at $p = 0.05$. For population analyses average firing rates were calculated for each neuron and the distributions of average firing rates were compared.

Results

Area LIP

In the rhesus monkey, where it has been most extensively characterized, the LIP occupies a small portion of the lateral bank of the intraparietal sulcus (Figure 1). Although a homologue of LIP is thought to exist in human parietal cortex, no consensus yet exists about its location and functional profile. Monkey LIP has extensive anatomical connections with two oculomotor structures, the frontal eye field (FEF) in the frontal lobe, and the superior colliculus (SC) in the midbrain, and weaker links with neighboring parietal areas that are related to skeletal (arm and hand) movements (Lewis and Van Essen, 2000a,b; Nakamura et al., 2001). In addition, LIP has bi-directional connections with extrastriate visual areas including motion selective areas – the middle temporal and middle superior temporal areas (MT and MST) – and shape and color selective areas including V4 and portions of the inferior temporal lobe (Lewis and Van Essen, 2000b). Finally, LIP is reciprocally connected with the posterior cingulate, a limbic area, and with the perirhinal and parahippocampal cortex, which comprise the gateway to the hippocampus (Blatt et al., 1990). Anatomically, therefore, LIP is well situated to receive visual, motor, motivational and cognitive information.

Consistent with their rich visual inputs, many LIP neurons have reliable visual responses and spatial selectivity. The fact that neurons have only weak selectivity in LIP for stimulus features (shape, color or pattern) suggests that visual inputs from extrastriate cortex provide convergent feature information. Receptive fields (RF) are typically contralateral to the recorded hemisphere and confined to a single quadrant (Platt and Glimcher, 1995; Ben Hamed et al., 2001). The LIP in each hemisphere comprises a complete representation of the contralateral field including the central, perifoveal region, with a rough topographic organization (Ben Hamed et al., 2001). RFs are retinotopic – that is, they are linked to the retina and move in space each time the eye moves, but neurons also receive extraretinal information regarding eye position and impending eye movements, which may be used to extract information in a more stable, world-referenced coordinate frames (Colby and Goldberg, 1999). Thus, LIP provides a topographic representation of contralateral visual space that combines retinal and extraretinal information.

In contrast with neurons in neighboring areas 7a or the ventral intraparietal area (VIP) that are best activated by large or full-field moving stimuli, LIP neurons respond exuberantly to small stationary objects flashed inside their RF. Visual onset responses occur with latencies as short as 40 ms and are notable for their precision and reliability (Bisley et al., 2004). However, despite their machine-like quality, these visual responses are not mere “sensory transients” but report the physical salience (conspicuity) of the stimuli eliciting them. Although neurons respond exuberantly to a stimulus that abruptly appears in the visual field, they have minimal or no responses if the same stimulus is stable in the environment and enters the RF by virtue of the monkeys’ eye movements (Gottlieb et al., 1998). This suggests that the visual signal in LIP reflects not the mere presence but the unusual conspicuity of a suddenly appearing object. Subsequent work verified this idea through direct measurements of the ability of flashed objects

to automatically draw attention. Bisley and Goldberg showed that the visual on response in LIP predicted the time course of rapid, covert attentional shifts toward a flashed distractor (Bisley and Goldberg, 2003). Balan and Gottlieb (Balan and Gottlieb, 2006) showed that neurons respond equally to several types of visual transients – including the appearance of a new object and an abrupt change in the color, position or luminance of an existing object, suggesting that they report an integrated measure of salience that is independent of specific stimulus properties (Figure 3a). These findings suggest that LIP visual responses reflect the end result of a rapid gating process, which filters out a large amount of information and selectively represents salient stimuli or events in the external world. The loci and mechanisms of this rapid gating are so far unknown.

In addition to their sensitivity to stimulus-driven salience, LIP neurons are also strongly sensitive to task-related factors. The earliest evidence for top-down modulations came from a study of Bushnell et al. who found that visual onset responses in LIP and in neighboring area 7a were enhanced if monkeys actively attended a stimulus in order to report its dimming relative to a passive viewing condition (Bushnell et al., 1981). Gottlieb et al. subsequently showed that top-down modulations were even more powerful in more naturalistic environments, in which stimuli were not intrinsically salient but had to be selected from a stable, multi-element display (Gottlieb et al., 1998). In this circumstance neurons responded minimally to a stable non-target object in their RF but developed robust responses if the object became relevant to the task. In other words, neurons clearly “pulled out” the task-relevant object from a complex visual scene.

The dynamics of the target selection signal in LIP correlate with behavioral measures of selection in tasks in which monkeys make eye movements toward the relevant target (Ipata et al., 2006; Thomas and Pare, 2007). A similar correlation is found if monkeys direct attention covertly, in the absence of overt gaze shifts. Balan et al. (Balan et al., 2008) trained monkeys to discriminate the orientation of a target presented among a variable number of distractors and report target orientation using a bar release, without shifting gaze to the target (Figure. 2a). The target was not physically conspicuous, and the task was challenging, with performance declining as the number of distractors increased, replicating the well-known set-size effect in human psychophysics (Carrasco and Yeshurun, 1998). LIP neurons developed a robust signal of target location, responding strongly if the target but not if a distractor was in their RF, even though monkeys maintained steady fixation (Figure 2b). In addition, their responses were affected by set size, declining as more distractors were added to the display (differences between red, blue and green curves in Figure 2b). The drop in firing rates was seen even before the active search phase, when monkeys were merely viewing a placeholder display with a variable number of elements, suggesting that it reflected competitive (mutually suppressive) visual interactions among the representations of individual elements. Set size continued to affect responses during the early part of the active search epoch (up to 200–300 ms after target onset), and the decline in target activity correlated with the set size effect in search performance (Figure 2c).

The findings suggest that LIP responses are limited by competitive visuo visual interactions among multiple stimuli, and these interactions can at least partly explain the capacity limitations measured behaviorally during visual search. Visuo-visual competition is also found in extrastriate visual areas that provide input to LIP, but, while in extrastriate cortex the competition is based on visual features (engage neurons with overlapping RF but disjoint feature preference) in LIP it is based on spatial location (engage neurons with no feature selectivity and disjoint RF). Thus competitive, mutually-suppressive interactions may limit processing at multiple levels of representation. Reversible inactivation of LIP with the GABA agonist muscimol (which transiently silences neuronal activity) produces deficits in finding targets in the contralateral visual field whose severity increases with set-size (Wardak et al.,

2002, 2004), consistent with the idea that LIP is critical for target selection and overcoming distractor competition in cluttered environments.

In addition to selecting among items of comparable salience, the brain must overcome interference from distractors that are more salient than the task-relevant target. In order to write this paper, for example, I must block out a persistent alarm in a nearby room and the sudden wail of an ambulance siren on the street. The brain's ability to control its own distractibility - the bottom up attentional weight of a task-irrelevant stimulus - is critical for maintaining focused, goal directed behavior. LIP neurons integrate information about bottom up and top-down priority, and reflect the weighting of salient stimuli according to their potential task relevance in a given context.

Balan and Gottlieb (Balan and Gottlieb, 2006) tested neural responses to salient visual transients presented during a covert visual search similar to that shown in Figure 2a. Monkeys performed the covert search task in Figure 2 with a constant number of distractors, but in this experiment, a transient visual perturbation (a change in one of the display elements) was presented for 200 ms before the search display. The perturbation was presented in two contexts. In one context the perturbation was informative for the search task always appearing at, and thus validly cueing, the target's location. In the other context the perturbation was uninformative, appearing at a location unrelated to the target. Because the two contexts were presented in separate blocks of trials monkeys could know ahead of time whether the visual transient was relevant or irrelevant to the task.

LIP neurons reflected behavioral context through a gain control mechanism. In the preparatory epoch prior to onset of the perturbation neurons showed a slight increase in their baseline firing rates in the relevant relative to the irrelevant condition (Figure 3b). This increase was remarkable because it was independent of spatial information and reflected a purely contextual, non-spatial signal. When the perturbation appeared neurons had robust on-responses to it, but these responses were stronger in the relevant than in the irrelevant context. Although the difference in firing rates between the two contexts was larger in the perturbation response than in the fixation epoch, the ratio of firing rates remained constant at 1.1 throughout the two epochs, suggesting that behavioral context changed a multiplicative gain of the neurons. The enhanced responses to the perturbation in the relevant context correlated with the facilitatory effect of the perturbation on reaction times (a decrease in reaction time; Figure 3c, left). The relatively lower responses in the irrelevant condition correlated with the distracting effect of the perturbation (an increase in reaction time; Figure 3c, right), which were overall weaker than the facilitatory effects. The findings show a complex integration of top-down and bottom up information in LIP. Although the priority map in LIP automatically registers salient changes in the world, the impact of such changes is subject to modification by behavioral context, or the subject's strategy, in this case through a multiplicative gain control mechanism.

Top down attention is task specific

The definition of bottom-up (or stimulus-driven) salience is relatively straightforward, as it is closely determined by the contrast between a stimulus and its environment in either a spatial dimension (e.g., a red apple among green apples) or in a temporal dimension (e.g., a sudden change in a stable object). The definition of "top down" attention, in contrast, is considerably more complex because it differs according to specific task demands. In the tasks discussed above monkeys directed attention with different goals in mind. In some cases attention was needed to select the target of an upcoming eye movement (Ipata et al., 2006; Thomas and Pare, 2007).. In others, attention was deployed covertly, in order to discriminate a peripheral visual target and emit a manual report (Figures 2 and 3). In yet other cases top-down attention incorporated prior knowledge about the structure of the task (Figure 3). The fact that LIP neurons encode attentional priority in all the circumstances described above suggests that the

brain's attentional architecture is to some degree centralized: at least some brain areas, such as LIP, compute attentional priority in a wide variety of behavioral circumstances.

This multifaceted engagement of LIP raises an interesting question: how does LIP learn how to select the object that is task relevant each specific behavioral context? Does this area receive information about all possible factors – sensory, motor, and cognitive – that could potentially influence attentional selection? If so, how does it integrate this information to deduce attentional priority? These questions present some of the key challenges for understanding the top-down control of attention. However, the existing body of evidence provides several clues. One clue is that neurons receive information about expected reward, which, by modulating visual and saccadic responses, could afford priority to stimuli that are more likely to lead to a desired outcome (Platt and Glimcher, 1999; Sugrue et al., 2004; Yang and Shadlen, 2007). Another clue is the emerging evidence that the priority map in LIP is not stereotyped but is modulated by decisional variables specific to various tasks. We suggest that both task-specific decisional inputs and reward signals act as teaching signals through which LIP in particular (and the brain in general) may identify and afford priority to task relevant sources of information.

The earliest decisional variable known to influence LIP was the planning of a saccadic eye movement. It is well known that LIP neurons show enhanced activity if monkeys plan saccades to visual stimuli, and, in some cases also respond for saccades toward blank spatial locations (Colby et al., 1996). Because a saccadic signal has properties that very closely resemble an attentional signal (enhanced activity at a particular spatial location) the two are often difficult, even impossible, to disambiguate in the neuronal firing rate. Recent findings from our laboratory, however, show that neurons also receive information about skeletal motor planning, and these signals are clearly distinct from the visuospatial response. As described above (Figure 2a) we used a covert visual search in which monkeys reported the orientation of a target by releasing a bar grasped with the right or left paw. While the primary signal in LIP during the search task was one of target location, with neurons responding more if the target than if a distractor appeared in the RF, many neurons (over 40% in two monkeys and over 70% in a third) were also strongly modulated by the bar release (Oristaglio et al., 2006). An example is shown in Figure 4. The neuron responded preferentially when the target than when a distractor was in the RF, but target responses were stronger if the monkey released the right bar than if she released the left bar. Control experiments established that these modulations reflected the active limb itself and not the visual properties of the target or the limb's position in space. Thus the neural visuospatial response in LIP was modulated by a skeletomotor, non-targeting manual response. A final study showed that decisional effects in LIP need not be linked to a specific motor act but can extend to more abstract judgments of stimulus category. Freedman and Assad used a delayed match-to sample task in which monkeys viewed a sample stimulus consisting of a field of dots drifting in one direction, and, after a delay, reported whether the direction of motion in a second, test, stimulus matched the category of the sample (Freedman and Assad, 2006). LIP neurons gave the expected on-response and sustained delay period activity if the task-relevant sample stimulus fell inside their RF, but in some neurons this response was modulated by target category, with different subsets of neurons showing preference for one of the two possible categories. Category preference changed when monkeys were retrained with a new category boundary, confirming that this selectivity reflected specific task demands. Thus, whether the operant requirement of the task is a saccade, a non-targeting manual response, or an abstract judgment of category membership, neurons combine a visuospatial signal of attentional priority with information about the operant response.

A new concept of attention

The complex combinatorial signals that bind visual-spatial and non-spatial information described in these studies were entirely unexpected on either experimental grounds, or based on theoretical understanding of the nature of attention. To understand this claim it is useful to outline two major concepts of attention that are popular in the research community.

One predominant view is that attention is a top-down control signal that affects visual (or more generally, sensory) representations so as to improve the information transmitted about attended relative to unattended objects (Reynolds and Chelazzi, 2004). This view is primarily concerned with the effects of attention on perception rather than with the processes underlying the attentional decision per se. An important implicit assumption of this view is that attentional effects are path-independent, i.e., they are independent of the context, or process by which attention has been allocated. Attention is thus regarded as a computation that is entirely unrelated to – and indeed, that must be carefully dissociated from – motor planning or decision making.

In contrast with this somewhat disembodied view, the premotor theory of attention makes the diametrically opposite claim, namely that that attention is a mere byproduct of motor planning (Rizzolatti et al., 1987; Moore et al., 2003). The premotor theory is strongly inspired by the neurophysiological evidence that structures implicated in attention overlap extensively with those controlling eye movements (Figure 1), and it specifically equates attention with saccade motor planning. In this view orienting attention is synonymous with planning an eye movement toward the attended target. While acknowledging that complex organisms can decouple attention and overt eye movements (i.e., they can direct attention covertly without shifting gaze or to distribute attention to more than one object at a time), proponents of the premotor theory maintain that attentional signals are essentially oculomotor signals that may or may not be implemented depending on downstream decision mechanisms.

The evidence emerging from LIP is consistent with the premotor theory in that it acknowledges the strong links between attention and rapid eye movements. However, this evidence considerably extends the premotor theory by showing that links between attention and decisions are much more varied and complex than previously believed. First, as describe above, activity related to visual selection is not exclusively linked with eye movement planning, but is also powerfully modulated by skeletomotor decisions and by abstract, categorical decisions. Second, the data clearly show that attentional signals represent distinct entities that cannot be reduced to mere motor planning signals. In the covert visual search task used in our lab (Figure 2 and Figure 4) LIP neurons encoded the location of the target even though that location was unrelated to the grasp release. In other words, the visuospatial response in LIP could not be interpreted as a motor planning signal, but could only be understood as an attentional response. Thus, although attention may have evolved from systems driving motor (perhaps specifically oculomotor) orienting, in high-order organisms such as primates it has gained its own neural machinery, which performs an internal selection process that is distinct from, although continuously informed by, decisional computations.

A computational model

How can we understand the hybrid signals in LIP in more formal, mechanistic terms? Answering this question will require much further research, but a particularly elegant beginning may be found in a computational model called Attention Gated Reinforcement Learning (AGREL), which elegantly combines the concepts of attention, decisions and learning (Roelfsema and van Ooyen, 2005). Although developed to emulate learning of pattern classification problems, the model shows properties that are strikingly similar to the findings described above in LIP.

AGREL consists of a three-layered neural network with an input, an intermediate and an output layer (Figure 5). Connections between layers are topographically organized, and include feedforward connections from the input to the middle layer and from the middle layer to the output layer, as well as feedback connections from the output layer to the middle layer. The network learns to solve various classification tasks - activate a single output in response to a given input pattern - using reinforcement learning. In the tradition of reinforcement learning, the model uses a reward prediction error (encoding the unexpected delivery of a reward) to gate plasticity. Learning is Hebbian, so that connections are strengthened among neurons that are active at the time of the reward. Also in keeping with reinforcement learning the reward signal is non-specific – i.e., it is broadcast to the entire network regardless of input-output patterns.

A key insight is that, although reward based learning does produce some desirable results, it is relatively inefficient because of the non-specific nature of the modifications. For example, imagine that inputs X_1 , X_2 and X_n are active at the time of a rewarded motor response Z_2 (Figure 5). Simple reward-driven plasticity will strengthen all synapses including those involving X_n , which were not specifically associated with the output. This non-specific strengthening of connections slows down the rate of learning. However, the model proposed that topographic feedback regarding the rewarded action can be used to improve the specificity of learning by assigning credit specifically to those units that led to that action. In the above example, competitive interactions in the output layer leads to selection of a single output, Z_2 . Feedback from the winning output in turn gates learning so that it occurs only at synapses directly involving this neuron. This feedback limits learning to synapses involving X_1 , X_2 and Y_1 , Y_2 , strengthening the pathway that gave rise to the rewarded action. The model thus elegantly predicts that middle layer units will have both sensory and motor properties and ascribes to the motor feedback a specific function in credit assignment and attentional learning.

While in its initial implementation the model replicated decisional effects in visual areas, the same architecture applies to a network in which the input represents a series of locations and the output a specific direction (i.e., an eye movement or attentional vector). This would directly capture the organization of the visual-attentional-oculomotor system into a cascade of retinotopically organized visual, attentional and saccade motor maps that are reciprocally and topographically interconnected. Because in natural behavior a salient object gives rise to an attentional response and to a saccade (i.e., activates an entire “column” in the network) this would naturally result in a middle (“attentional”) layer that shows converging visual and motor responses for a specific part of the visual field. An important challenge will be to extend the model to include both pattern and spatial information and test whether it reproduces the combinatorial coding of spatial and non-spatial signals experimentally encountered in LIP.

Conclusions

“Every one knows what attention is,” famously wrote William James in *Principles of Psychology* in 1890. “It is the taking possession by the mind, in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought. Focalization, concentration of consciousness are of its essence.” Investigations inside the brain’s black box show that the computations underlying attention are not quite as obvious as our common intuition suggests. However, the complex signals underlying attention appear to have some identifiable ingredients: the stimulus-driven salience of an external object, information about the type of decision related to that object, and information about expected outcome. Discovering how these signals are combined, and how they are molded together during learning of new behaviors, will be among the most rewarding aspects of upcoming research.

References

- Balan PF, Gottlieb J. Integration of exogenous input into a dynamic salience map revealed by perturbing attention. *J Neurosci* 2006;26:9239–9249. [PubMed: 16957080]
- Balan PF, Oristaglio J, Schneider DM, Gottlieb J. Neuronal correlates of the set-size effect in monkey lateral intraparietal area. *PLoS Biology*. 2008in press
- Ben Hamed S, Duhamel JR, Bremmer F, Graf W. Representation of the visual field in the lateral intraparietal area of macaque monkeys: a quantitative receptive field analysis. *Exp Brain Res* 2001;140:127–144. [PubMed: 11521146]
- Bisley JW, Goldberg ME. Neuronal activity in the lateral intraparietal area and spatial attention. *Science* 2003;299:81–86. [PubMed: 12511644]
- Bisley JW, Krishna BS, Goldberg ME. A rapid and precise on-response in posterior parietal cortex. *J Neurosci* 2004;24:1833–1838. [PubMed: 14985423]
- Blatt GJ, Andersen RA, Stoner GR. Visual receptive field organization and cortico-cortical connections of the lateral intraparietal area (area LIP) in the macaque. *J Comp Neurol* 1990;299:421–445. [PubMed: 2243159]
- Bushnell MC, Goldberg ME, Robinson DL. Behavioral enhancement of visual responses in monkey cerebral cortex: I. Modulation in posterior parietal cortex related to selective visual attention. *J Neurophysiol* 1981;46:755–772. [PubMed: 7288463]
- Carrasco M, Yeshurun Y. The contribution of covert attention to the set-size and eccentricity effects in visual search. *J Exp Psychol Hum Percept Perform* 1998;24:673–692. [PubMed: 9554103]
- Colby CL, Goldberg ME. Space and attention in parietal cortex. *Annu Rev Neurosci* 1999;23:319–349. [PubMed: 10202542]
- Colby CL, Duhamel J-R, Goldberg ME. Visual, presaccadic and cognitive activation of single neurons in monkey lateral intraparietal area. *J Neurophysiol* 1996;76:2841–2852. [PubMed: 8930237]
- Freedman DJ, Assad JA. Experience-dependent representation of visual categories in parietal cortex. *Nature* 2006;443:85–88. [PubMed: 16936716]
- Gottlieb J, Kusunoki M, Goldberg ME. The representation of visual salience in monkey parietal cortex. *Nature* 1998;391:481–484. [PubMed: 9461214]
- Ipata AE, Gee AL, Goldberg ME, Bisley JW. Activity in the lateral intraparietal area predicts the goal and latency of saccades in a free-viewing visual search task. *J Neurosci* 2006;26:3656–3661. [PubMed: 16597719]
- Lewis JW, Van Essen DC. Mapping of architectonic subdivisions in the macaque monkey, with emphasis on parieto-occipital cortex. *J Comp Neurol* 2000a;428:79–111. [PubMed: 11058226]
- Lewis JW, Van Essen DC. Corticocortical connections of visual, sensorimotor, and multimodal processing areas in the parietal lobe of the macaque monkey. *J Comp Neurol* 2000b;428:112–137. [PubMed: 11058227]
- Moore T, Armstrong KM, Fallah M. Visuomotor origins of covert spatial attention. *Neuron* 2003;40:671–683. [PubMed: 14622573]
- Nakamura H, Kuroda T, Wakita M, Kusunoki M, Kato A, Mikami A, Sakata H, Itoh K. From three-dimensional space vision to prehensile hand movements: the lateral intraparietal area links the area V3A and the anterior intraparietal area in macaques. *J Neurosci* 2001;21:8174–8187. [PubMed: 11588190]
- Oristaglio J, Schneider DM, Balan PF, Gottlieb J. Integration of visuospatial and effector information during symbolically cued limb movements in monkey lateral intraparietal area. *J Neurosci* 2006;26:8310–8319. [PubMed: 16899726]
- Platt ML, Glimcher PW. Response fields of neurons in area LIP during performance of delayed saccade and selection tasks. *Soc Neurosci Abstr* 1995;21:1196.
- Platt ML, Glimcher PW. Neural correlates of decision variables in parietal cortex. *Nature* 1999;400:233–238. [PubMed: 10421364]
- Reynolds JH, Chelazzi L. Attentional modulation of visual processing. *Annu Rev Neurosci* 2004;27:611–647. [PubMed: 15217345]

- Rizzolatti G, Riggio L, Dascola I, Umiltá C. Reorienting attention across the horizontal and vertical meridians: Evidence in favor of a premotor theory of attention. *Neuropsychologia* 1987;25:31–40. [PubMed: 3574648]
- Roelfsema PR, van Ooyen A. Attention-gated reinforcement learning of internal representations for classification. *Neural Comput* 2005;17:2176–2214. [PubMed: 16105222]
- Sugrue LP, Corrado GS, Newsome WT. Matching behavior and the representation of value in the parietal cortex. *Science* 2004;304:1782–1787. [PubMed: 15205529]
- Thomas NW, Pare M. Temporal processing of saccade targets in parietal cortex area LIP during visual search. *J Neurophysiol* 2007;97:942–947. [PubMed: 17079346]
- Wardak C, Olivier E, Duhamel JR. Saccadic target selection deficits after lateral intraparietal area inactivation in monkeys. *J Neurosci* 2002;22:9877–9884. [PubMed: 12427844]
- Wardak C, Olivier E, Duhamel JR. A deficit in covert attention after parietal cortex inactivation in the monkey. *Neuron* 2004;42:501–508. [PubMed: 15134645]
- Yang T, Shadlen MN. Probabilistic reasoning by neurons. *Nature* 2007;447:1075–1080. [PubMed: 17546027]




Figure 1. Lateral view of the macaque brain

showing the location of area LIP and the principal areas, on the lateral aspect of the cortical hemisphere, with which LIP is connected (in red). Abbreviations: **cal**, calcarine sulcus; **lun**, lunate sulcus; **sts**, superior temporal sulcus; **lf**, lateral fissure; **cs**, central sulcus; **as**, arcuate sulcus; **ps**, principal sulcus; **MT**, middle temporal area; **MST**, middle superior temporal area; **FEF**, frontal eye field; **IT**, inferior temporal cortex; **SC**, superior colliculus. Numbers represent cortical areas.

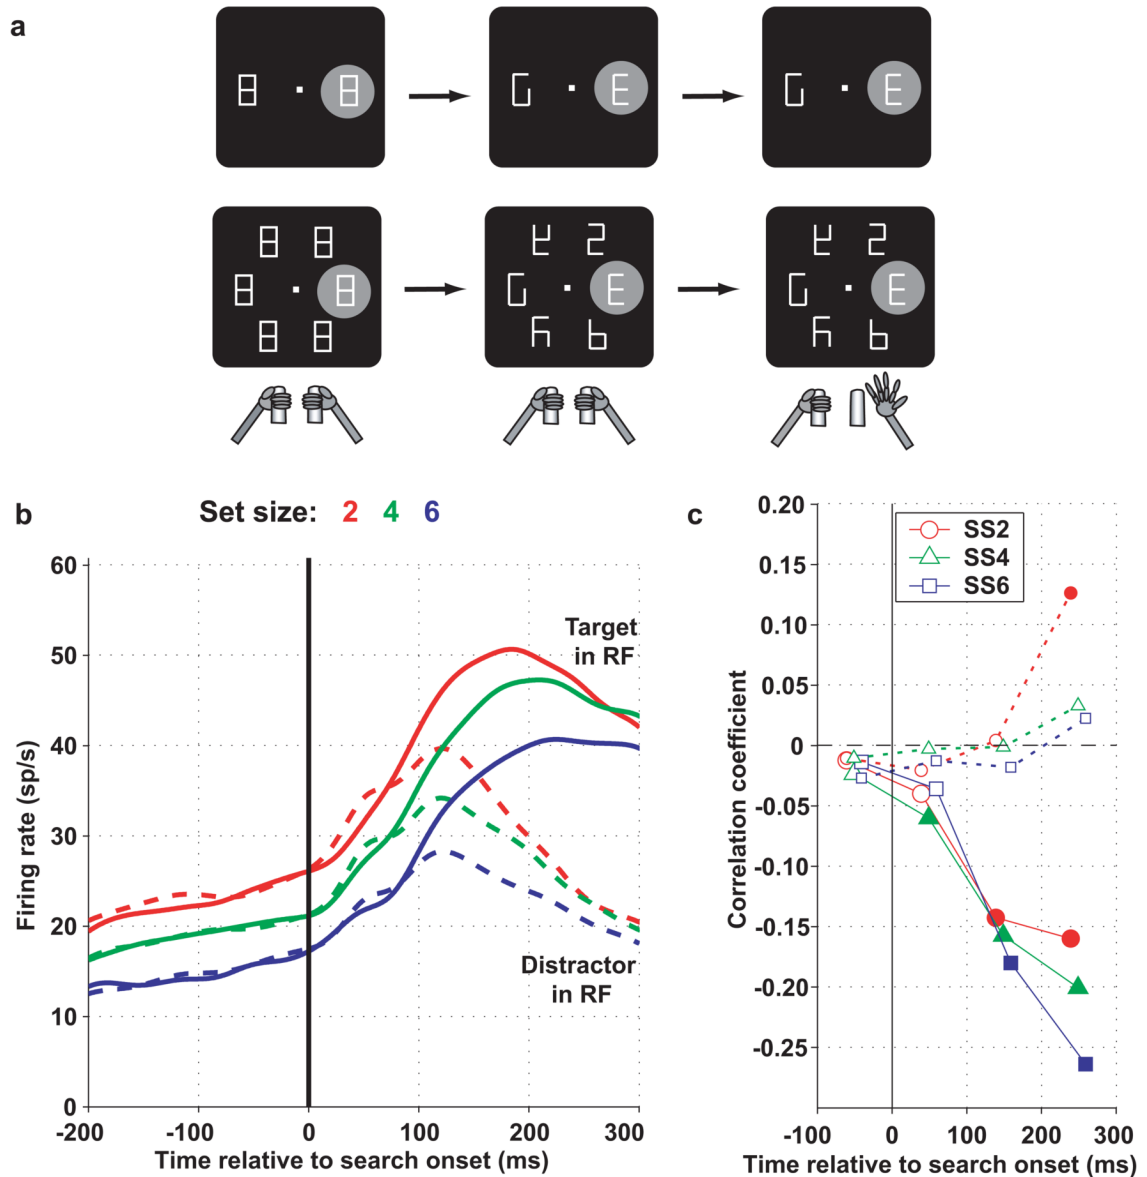


Figure 2. LIP neurons encode the locus of covert attention a. Behavioral task

A display containing 2 or 6 figure-8 placeholders remained stably on the screen and, to begin a trial, monkeys fixated a central spot and grabbed two bars at stomach level (left panel). When the monkey achieved fixation the RF of the neuron under investigation fell upon one of the placeholders. The RF is shown as the gray shaded area, for a putative neuron with RF to the right of the fovea. However in practice neural RF could cover any portion of the visual screen and the display was scaled and rotated to match the coordinates of the RF. Approximately 500 ms after fixation achievement two line segments were removed from each placeholder, revealing one cue, the letter “E”, and several unique distractors (middle). Unpredictably and with uniform probabilities, the “E” appeared at any display location and could be forward-facing (as shown) or backward-facing (not shown). Monkeys received a juice reward if they indicated the orientation of the E by releasing the right bar if the E was forward facing or the left bar if it was backward-facing. Trials ended with extinction of the fixation point and restoration of the placeholder display. **b.** Population responses on trials in which the target

(solid) or a distractor (dashed) were in the RF at set sizes 2, 4 and 6. Spike density histograms were derived by convolving individual spike times with a Gaussian kernel with standard deviation of 15 ms and averaging the convolved traces across all neurons ($n = 50$ from 2 monkeys). Responses are aligned on search display onset (presentation of the search display, time 0). Prior to search onset the monkeys viewed the stable placeholder display with a variable number of elements, and firing rates declined as a function of set size. During search firing rates were higher if the target than if a distractor was in the RF, and the effect of set size diminished gradually by 300 ms after display onset. **c.** Correlation coefficients between firing rates and reaction times on a trial by trial basis. Each point shows the correlation coefficient between reaction time and firing rate in a 100 ms bin, with significant values ($p < 0.05$) denoted by filled symbols. Responses to the target in the RF were significantly, negatively correlated with reaction times (higher responses were associated with shorter reaction times). Responses to a distractor in the RF had much weaker correlations. Adapted, with permission, from (Balan et al., 2008).

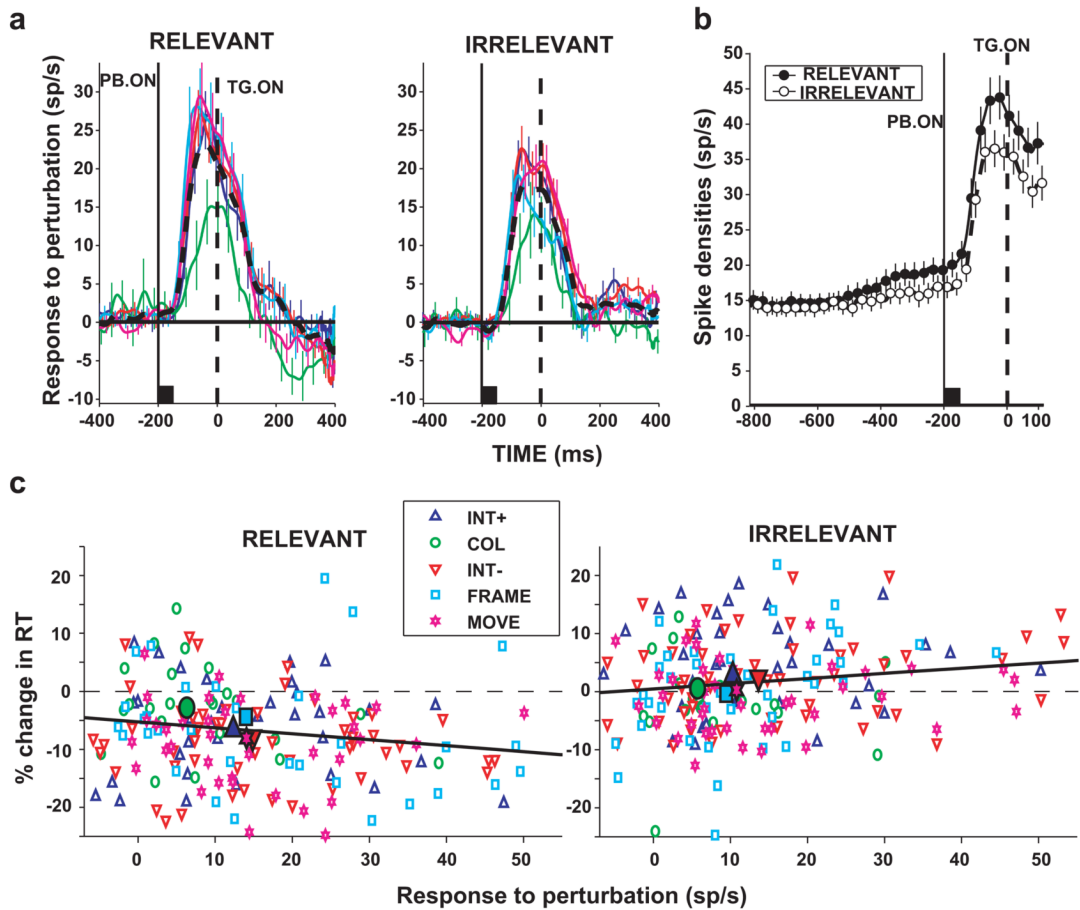


Figure 3. LIP neurons integrate stimulus and task information

Monkeys performed a search task similar to that shown in Figure 2a, but, 200 ms prior to search onset, were shown a brief visual perturbation, a 50 ms change in one of the display elements.

a. Population responses to transient visual perturbations consisting of a brief increase or decrease in luminance, a change of color, a sudden motion, or appearance of a FRAME (legend in panel c). The time of search onset is indicated by the dashed line (TG. ON) and perturbation onset by a solid line (PB. ON). Neurons responded similarly to all perturbation types, but responses were overall stronger when the perturbation was relevant than when it was irrelevant to the task. **b.** Population responses during the fixation epoch and the perturbation epoch. Before perturbation onset activity increased slightly in the relevant relative to the irrelevant condition. This enhancement was more marked when the perturbation appeared (PB. ON, -200 ms), reflecting a multiplicative gain. Symbols show average and standard error of the firing rate. **c.** Relationship between responses to the perturbation and the perturbation effect on reaction times (difference between reaction time on perturbation and no-perturbation trials). In the relevant context the perturbation produced shortened reaction times (facilitatory effect) and in the irrelevant context it lengthened reaction times (distracting effect). Adapted, with permission, from (Balan and Gottlieb, 2006).

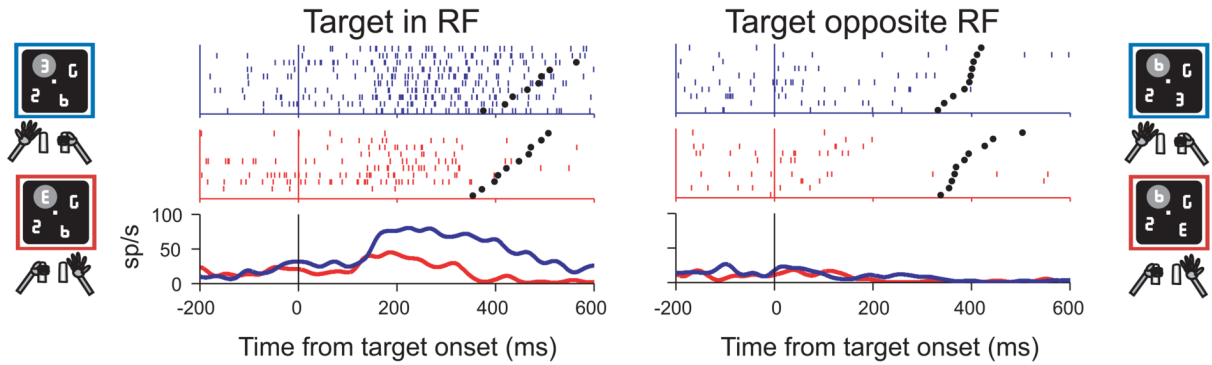


Figure 4. Spatial responses to the search target are modulated by limb motor planning
 Representative limb effect in one neuron. Neural activity is shown for trials in which the target was in the center of the RF or at the opposite location and was either facing to the right (red) or to the left (blue). In the raster plots each line is one trial, each tick represents the time of one action potential relative to cue presentation. Spike density histograms were derived by convolving individual spike times with a Gaussian kernel with standard deviation of 15 ms. Activity is aligned on cue presentation (time 0) and black dots represent bar release. Only correct trials are shown, ordered offline according to reaction time. Cartoons indicate the location of the cue, the RF (gray oval) and the manual response. The neuron responded best when the cue was in its RF (upper left quadrant) and the monkey released the left bar. Adapted, with permission, from (Oristaglio et al., 2006)

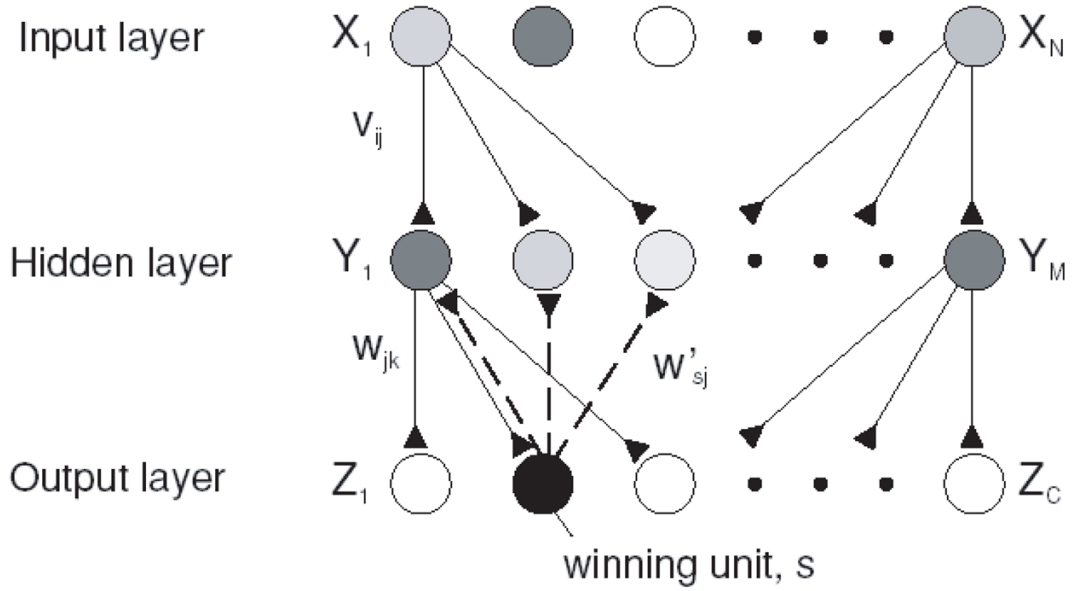


Figure 5. Three layered network used to perform classification tasks

The task of the network is to activate a single output unit that encodes the class of a stimulus pattern. There are N units in the input layer, M in the second (hidden) layer and C in the output layer. Connections v_{ij} propagate activity from the input to the hidden layer, and connections w_{jk} propagate activity from the hidden to the output layer. The winning unit, s , feeds its activity back to the hidden layer through connections w'_{sj} (dashed lines). Adapted, with permission, from (Roelfsema and van Ooyen, 2005).