

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

Jacqueline Gottlieb^{a,b,*}, Puiu Balan^a, Jeff Oristaglio^a, Mototaka Suzuki^a

^a Department of Neuroscience, Columbia University, 1051 Riverside Drive, Kolb Research Annex, Room 569, New York, NY 10032, USA

^b Department of Psychiatry, Columbia University, 1051 Riverside Drive, Kolb Research Annex, Room 569, New York, NY 10032, USA

ARTICLE INFO

Article history:

Received 3 July 2008

Revised 16 September 2008

Accepted 16 September 2008

Available online 8 November 2008

Keywords:

Attention
Parietal cortex
Priority map
Learning

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 of various stimuli based on the task-specific associations between these stimuli, the required action and expected outcome.

© 2008 Elsevier Inc. All rights reserved.

1. Introduction

Judicious selection is at the heart of goal directed behavior. To select appropriately in a complex environment an intelligent agent (be it a person, a rat or a monkey), must solve two problems. First, the agent must identify the most relevant source of information from among many potential alternatives. Second, it must be able to focus on the relevant source and block out irrelevant distractions. In both intuitive and scientific terms, we think of the cognitive operations that allow such adaptive selection as falling into the broad realm of “attention”.

It is often claimed that attention is necessary for overcoming capacity limitations inherent in neural processing. Because the brain is “bombarded” with more sensory information than it can process in depth, the argument goes, attention is needed to prioritize and limit the amount of information that reaches higher processing stages at any one time. However, the need for selection remains even in simple environments that do not seriously tax capacity limitations. Even in such environments, we must decide which objects are helpful to us and which are not. 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, actions and outcomes. It follows that attention 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 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. Investigations into LIP began with the somewhat naïve view that an area controlling attention must simply represent objects or locations that are attended at a given moment and respond relatively weakly to distractors. To a first approximation, this is indeed what is found in LIP. However, recent evidence shows that this “priority map” is more complex and in particular that it takes on a wide range of task-specific properties – i.e., it appears to be plastic and adaptable to task demands. The specific significance of these modulations is not fully understood. However, we suggest that these modulations represent teaching signals through which the brain learns to assign attentional priority to various stimuli based on their significance for a specific action or a specific outcome. We end by describing a computational model (Roelfsema & van Ooyen, 2005) that may be a good starting point for formalizing inquiry into the links between learning and attention.

2. Methods

2.1. General methods and behavioral tasks

Data were collected with standard behavioral and neurophysiological techniques as described previously (Balan & Gottlieb, 2006;

* Corresponding author. Address: Department of Neuroscience, Columbia University, 1051 Riverside Drive, Kolb Research Annex, Room 569, New York, NY 10032, USA. Fax: +1 212 543 5816.

E-mail address: jg2141@columbia.edu (J. Gottlieb).

Oristaglio, Schneider, Balan, & Gottlieb, 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.

2.2. 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, $p < .05$) 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).

2.3. Covert search task

The basic variant of the covert search task (Fig. 2a) was tested with display size of four 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 (Fig. 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 randomly 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.

2.4. 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 $p = .05$. For population analyses average firing rates were calculated for each neuron and the distributions of average firing rates were compared.

3. Results

3.1. 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 (Fig. 1). Although a homologue of LIP is thought to exist in human parietal cortex, no consensus yet exists about its location and functional profile. Anatomically, LIP is well situated to receive visual, motor, motivational and cognitive information. It has extensive anatomical connections with an oculomotor structure, the frontal eye field (FEF) in the frontal lobe, and weaker links with neighboring parietal areas that are related to skeletal (arm and hand) movements (Lewis & Van Essen, 2000a, 2000b; Nakamura et al., 2001). It projects to the superior colliculus, a subcortical oculomotor area. In addition, LIP has bi-directional connections with extrastriate visual areas including motion selective areas in the dorsal stream and shape and color selective areas in the ventral stream (Lewis & 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, Andersen, & Stoner, 1990).

Consistent with their rich visual inputs, many LIP neurons have visual responses and spatial receptive fields (RF), which are typically contralateral to the recorded hemisphere and confined to a single quadrant (Ben Hamed, Duhamel, Bremmer, & Graf, 2001). Thus, the LIP in each hemisphere comprises a complete representation of the contralateral field with a rough topographic organization, that include a representation of the perifoveal region, (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. However, 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 & Goldberg, 1999). Thus, the topographic representation in LIP is informed of the organism's position and is suitable for generating motor commands directed toward specific spatial locations.

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 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, Krishna, & Goldberg, 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. This has been shown in a task in which an array of stimuli remained stable on the screen, and monkeys made eye movements that brought these stimuli into the RF (Gottlieb, Kusunoki, & Goldberg, 1998). Neurons had little response to the stable, stimuli if these were not relevant to the task. Thus, complex, natural scenes evoke relatively little activity in LIP, but neurons selectively respond to intrinsically salient, flashed objects.

LIP responses to flashed stimuli correlate with automatic shifts of attention to salient objects. 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 & Goldberg, 2003). Balan and Gottlieb (Balan & Gottlieb, 2006) showed that neurons respond equally to visual transients defined by different physical characteristics, such as an abrupt change in color, position or luminance (Fig. 3a).

In addition to their sensitivity to stimulus-driven salience, 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 neighbor-

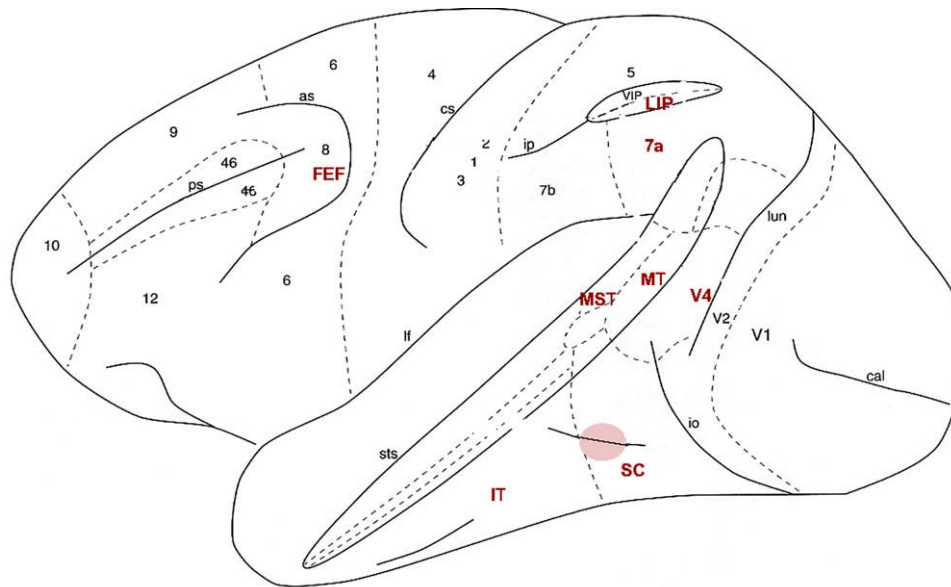


Fig. 1. Lateral view of the macaque brain showing the location of area LIP and the principal areas, on the lateral aspect of the cortical hemisphere, with which LIP is connected (in red). Abbreviations: cal, calcarine sulcus; lun, lunate sulcus; sts, superior temporal sulcus; lf, lateral fissure; cs, 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. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this paper.)

ing area 7a were enhanced if monkeys actively attended a stimulus in order to report its dimming, relative to a passive viewing condition (Bushnell, Goldberg, & Robinson, 1981). Gottlieb et al. subsequently showed that top-down modulations were even more powerful in more naturalistic environments, in which stimuli were not flashed but had to be selected from a stable, multi-element display (Gottlieb et al., 1998). In this circumstance neurons responded minimally to a task-irrelevant stable object but developed robust responses if the object became relevant to the task. In other words, neurons “pulled out” the task-relevant objects from complex scenes.

As for the salience-driven responses, these top-down target selection signals correlate with behavioral performance (Ipata, Gee, Goldberg, & Bisley, 2006; Thomas & Pare, 2007). A recent study demonstrates these correlations in a task in which monkeys directed attention covertly to a target during visual search (Balan, Oristaglio, Schneider, & Gottlieb, 2008). Balan, Oristaglio, Schneider, and Gottlieb (2008) trained monkeys to discriminate the orientation of a target (a letter “E”) that was presented among letter-like distractors (Fig. 2a). The target appeared at an unpredictable location in the display and, in order to identify it the monkeys had first to find it. However, monkeys were not allowed to make eye movements to the search display, so that target selection and discrimination had to be done covertly. Even though eye movements were not allowed, LIP neurons developed a robust signal of target location, responding strongly if the target but not if a distractor was in their RF (Fig. 2b). This demonstrates that neurons are as effective in “pulling out” behavioral targets for covert operations as they are when frank physical orienting is required. Moreover, the properties of the target response in LIP corresponded with the monkeys’ performance. Because the target was not physically conspicuous, performance (as measured through reaction time and accuracy) declined as the number of distractors increased. While this so called set-size effect is well-known from decades of research in human psychophysics, its neural substrates had remained unknown (Carrasco & Yeshurun, 1998). However, in LIP target-related responses declined as more distractors were added

to the display, showing a neural set-size effect that corresponded to the behavioral set-size effect (differences between red, blue and green curves in Fig. 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, and declined during active search. This suggests that the neural set-size effect reflects competitive interactions among the representation of multiple display elements, but these interactions are resolved during active search, as the brain finds the target. Thus, LIP activity reflects the underlying dynamics and difficulty of the process of target selection.

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 contralesional visual field whose severity increases with set-size (Wardak, Olivier, & Duhamel, 2002, 2004), consistent with the idea that LIP is critical for target selection and overcoming distractor competition in cluttered environments.

Since LIP neurons combine information about physical salience with that about task-relevance, one must ask how they balance the potentially competing demands of sensory and internal cues that attract attention. In natural behavior, we need to afford more or less attention to a salient object depending on behavioral context. For example, if we are crossing the street we would do well to attend to the loud wail of an ambulance; if we are trying to write a paper in the office, however, it would be best to block out that wail as much as possible. In order to deploy attention effectively the brain must balance the competing demands of bottom-up (sensory) and top-down (voluntary) cues according to behavioral context. Will LIP neurons reflect this balance?

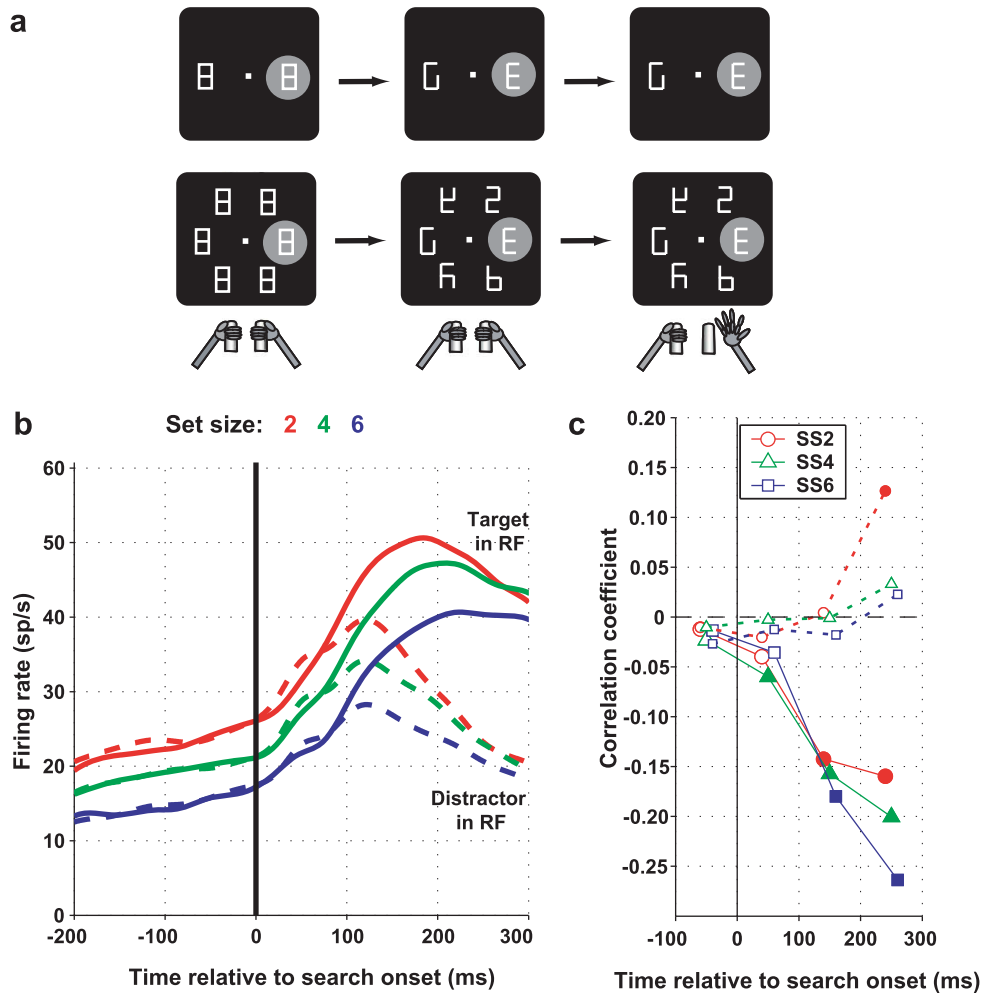


Fig. 2. LIP neurons encode the locus of covert attention. (a) Behavioral task. A display containing several figure-8 placeholders remained stably on the screen and, to begin a trial, monkeys fixated a central spot and grabbed two bars at stomach level (left panel). When the monkey achieved fixation the 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 ($p < .05$) 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).

A recent study shows that they do. Balan and Gottlieb (2006) tested neural responses to salient visual transients presented during a covert visual search similar to that shown in Fig. 2a. Monkeys performed the search task in Fig. 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, because it always appeared at, and thus validly indicated the location at which the target will appear. 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 knew 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 before presentation of any stimulus, neurons showed a slight increase in their baseline firing rates in the relevant relative to the irrelevant context (Fig. 3b). This increase was remarkable because it indicated a non-sensory, preparatory signal. When the perturbation appeared, neurons responded to it, and these responses were stronger in the relevant than in the irrelevant context. The ratio of firing rates between the relevant and irrelevant contexts remained constant at 1.1 throughout the preparatory and perturbation epochs, suggesting that behavioral context had a multiplicative effect. Neural responses to the perturbation correlated with the effect of the perturbation on reaction times. In the irrelevant condition the perturbation produced an increase in reaction time (a distracting effect) which was higher if the perturbation-evoked response

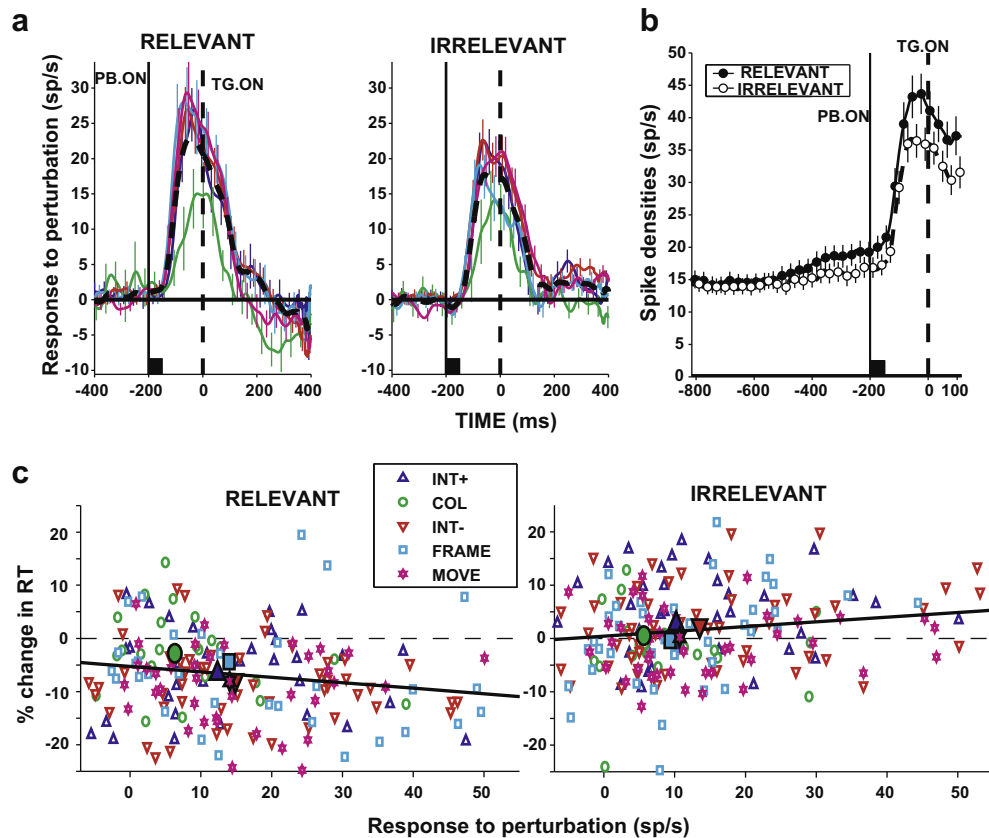


Fig. 3. LIP neurons integrate stimulus and task information. Monkeys performed a search task similar to that shown in Fig. 2a, but, 200 ms prior to search onset, were shown a brief visual perturbation, a 50 ms change in one of the display elements. (a) Population responses to transient visual perturbations consisting of a brief increase or decrease in luminance, a change of color, a sudden motion, or appearance of a FRAME (legend in 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).

was higher (Fig. 3c, right). In the relevant condition the perturbation produced a decrease in reaction times (a facilitatory effect), which also increased as the perturbation-evoked response increased. Thus, LIP neurons integrate top-down and bottom-up information in context-dependent manner. Although the priority map in LIP automatically registers salient changes in the world, the impact of such changes depends on the subject's prior knowledge and behavioral strategy.

3.2. 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, and differs according to specific task demands. In the experiments mentioned 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 & Pare, 2007). In others, attention was deployed covertly, in order to discriminate a visual target and emit a manual report (Figs. 2 and 3). In yet other cases top-down attention incorporated prior knowledge about the structure of the task (Fig. 3). The fact that LIP neurons encode attentional priority in all the circumstances de-

scribed above suggests that the brain's attentional architecture is to some degree centralized: at least some brain areas compute attentional priority in a wide variety of behavioral circumstances.

This observation raises an interesting question: how does LIP learn how to select the relevant object in each 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? The answers to these questions are not yet known, but, the existing evidence provides a few tantalizing 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 & Glimcher, 1999; Sugrue, Corrado, & Newsome, 2004). Another clue is the emerging evidence that the priority map in LIP 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.

"Decisional factors" represent the final goal, or decision that is involved in a behavioral task. An obvious decisional variable is the motor response that must be emitted by the subject. It is well known that LIP is sensitive to one type of movement, rapid eye movements (saccades) toward attended targets. While we have seen above that neurons encode covert attention even when there

is no motor orienting (Figs. 2 and 3), many neurons show enhanced activity if monkeys plan actual saccades to attended stimuli, and some respond for saccades toward blank spatial locations (Colby, Duhamel, & Goldberg, 1996). Because saccades require a topographic representation of space similar to that required for spatial attention, it may seem unsurprising that a common population of neurons is activated for attention and saccades. However, recent findings from our laboratory, show that neurons also receive information about motor decisions that are not coded in visuospatial coordinates. In the covert visual search task described above (Fig. 2a) monkeys reported the orientation of the “E” target with a non-spatial manual report, by releasing a bar grasped with the right or left paw. As we described above, the primary signal in LIP was one of target location, with neurons responding more if the target than if a distractor appeared in the RF. However, in many neurons this signal was modulated by the bar release (Oristaglio et al., 2006). An example is shown in Fig. 4. This neuron responded more strongly to the target if the monkey released the left bar than if she released the right bar. If a distractor was in the RF the neuron did not respond for either bar release. Thus, the neuron did not encode the motor act itself. However, it incorporated information about the motor act in its attentional response. Control experiments established that these modulations reflected the active limb (effector) itself and not the visual properties of the target or the limb’s position in space. Another 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 found that, if monkeys report the category of a moving stimulus, neurons respond strongly to the moving stimulus in their RF but these responses were modulated by stimulus category, with different subsets of neurons showing preference for one of the two possible categories (Freedman & Assad, 2006). Thus, the visual salience map in LIP is not purely visual but is shaped by multiple decisional factors. These range from factors such as eye movements that are closely related with attention, to non-spatial motor act such as a manual grasp, to abstract cognitive factors such as categorization.

3.3. A new concept of attention

The complex binding of visual–spatial and non-spatial information in LIP was entirely unexpected based on the intuitive concept of attention that motivated these studies in the first place. To understand this claim it is useful to outline two major concepts of attention that are popular in the research community.

One 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 the attended object (Reynolds & Chelazzi, 2004). This view is primarily concerned with the effects of attention on perception, and does not say much about the process by which the brain decides where to attend. This view makes the implicit assumption that attentional effects are independent of the context, or process by which attention has been allocated. In this view therefore, attention is a relatively disembodied entity – a computation that is entirely unrelated to, and thus must be carefully dissociated from, motor planning or decision making.

In stark contrast with this view, the premotor theory of attention makes the claim, that attention is a mere byproduct of motor planning (Moore, Armstrong, & Fallah, 2003; Rizzolatti, Riggio, Dascola, & Umiltà, 1987). Inspired by neurophysiological evidence that the neural substrates of attention overlap with those controlling eye movements, the premotor theory specifically equates attention with saccade motor planning. Proponents of the premotor theory maintain that attentional signals are essentially oculomotor signals, and that these signals may or may not be implemented (an actual eye movement may or may not occur), depending on downstream decision mechanisms.

The evidence emerging from LIP suggests a synthesis of these two views. This evidence acknowledges is consistent with the premotor theory in that it acknowledges the strong links between attention and motor planning. However, the evidence shows that these links reflect a broad space of decisional variables beyond eye movements (i.e., skeletal movements and abstract decisions), and also that attention is not synonymous with motor planning. In the covert visual search task used in our lab (Fig. 4) LIP neurons did not consistently encode the grasp release. Instead, they encoded target location (even though target location was not statistically related with the manual report). It is entirely possible that attention evolved from systems that drove motor orienting, in particular the superior colliculus. However, in high-order organisms such as primates attention gained its own neural machinery. This machinery now performs an internal selection process that is distinct from, although continuously informed by, overt actions.

3.4. 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 good 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 & van Ooyen, 2005). Although developed to emulate learning of pattern classification problems,

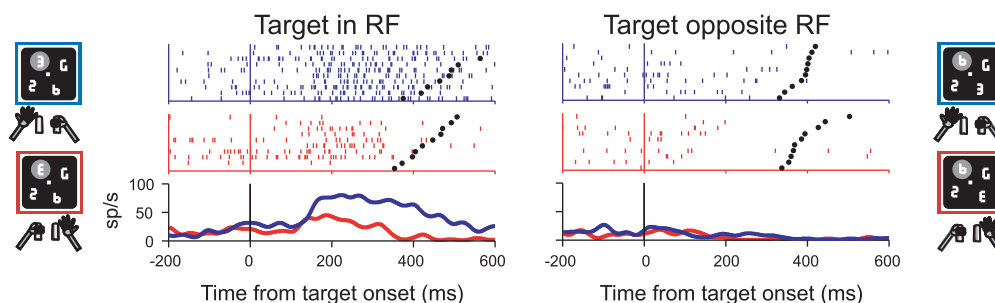


Fig. 4. Spatial responses to the search target are modulated by limb motor planning. Representative limb effect in one neuron. Neural activity is shown for trials in which the target was in the center of the RF or at the opposite location and was either facing to the right (red) or to the left (blue). 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). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this paper.)

the model shows properties that are strikingly similar to the findings described in LIP.

AGREL consists of a three-layered neural network with an input, an intermediate layer and an output layer (Fig. 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 (generated by 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 can produce results, it is relatively inefficient because it is non-specific. For example, imagine that inputs X_1 , X_s and X_n are active at the time of a rewarded motor response Z_s (Fig. 5). Simple reward-driven plasticity will strengthen all three inputs. However, only X_1 and X_s directly drove the motor output while X_n was coincidentally active. For instance, this input may represent a salient object that happens to attract attention but is not relevant for the motor response. Thus, simple reward based learning can slow learning by inappropriately assigning credit to external inputs. AGREL proposes that topographic feedback regarding the rewarded action can be used to improve learning, by modifying specifically those inputs that led to the action. In the above example, competitive interactions in the output layer lead to selection of a single output, Z_s . Feedback from the winning output in turn gates learning so that it occurs only at synapses directly involving this neuron, and thus selectively strengthens the pathway that gave rise to the rewarded action (in this example, inputs from X_1 , X_s and Y_1 , Y_s). 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 – deciding where to attend given specific task requirements.

In its initial implementation the model consists of a single input layer (X), coding a single visual dimension. An important challenge will be to extend the model to include more than one input dimension, for example, information about visual features and spatial




Fig. 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).

information, coded in separate neural maps (layers). One can then observe the combinatorial coding that emerges in intermediate layers as the network learns specific tasks.

4. 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.” This is indeed a compelling intuition, and one which guided the first forays of neurophysiologists into the black box that controls attention. These forays to some extent confirmed our intuition. Neurons in LIP, as well as those in the superior colliculus and the frontal eye field (and possibly in other areas), do seem to track a continuously moving mental spotlight, the “taking possession by the mind” of different objects at different points in time. However, these neurons also carry signals that are not quite as intuitive as suggested by William James, and whose role in attention is not entirely understood. Testing the idea that these signals are ingredients of the complex process by which the brain decides where to attend promises to be among the most rewarding aspects of upcoming research.

References

- Balan, P. F., & Gottlieb, J. (2006). Integration of exogenous input into a dynamic salience map revealed by perturbing attention. *The Journal of Neuroscience*, 26, 9239–9249.
- Balan, P. F., Oristaglio, J., Schneider, D. M., & Gottlieb, J. (2008). Neuronal correlates of the set-size effect in monkey lateral intraparietal area. *PLoS Biology*, 6, e158.
- Ben Hamed, S., Duhamel, J. R., Bremmer, F., & Graf, W. (2001). Representation of the visual field in the lateral intraparietal area of macaque monkeys: A quantitative receptive field analysis. *Experimental Brain Research*, 140, 127–144.
- Bisley, J. W., & Goldberg, M. E. (2003). Neuronal activity in the lateral intraparietal area and spatial attention. *Science*, 299, 81–86.
- Bisley, J. W., Krishna, B. S., & Goldberg, M. E. (2004). A rapid and precise on-response in posterior parietal cortex. *The Journal of Neuroscience*, 24, 1833–1838.
- Blatt, G. J., Andersen, R. A., & Stoner, G. R. (1990). Visual receptive field organization and cortico-cortical connections of the lateral intraparietal area (area LIP) in the macaque. *The Journal of Comparative Neurology*, 299, 421–445.
- Bushnell, M. C., Goldberg, M. E., & Robinson, D. L. (1981). Behavioral enhancement of visual responses in monkey cerebral cortex: I. Modulation in posterior parietal cortex related to selective visual attention. *Journal of Neurophysiology*, 46, 755–772.
- Carrasco, M., & Yeshurun, Y. (1998). The contribution of covert attention to the set-size and eccentricity effects in visual search. *Journal of Experimental Psychology: Human Perception and Performance*, 24, 673–692.
- Colby, C. L., Duhamel, J.-R., & Goldberg, M. E. (1996). Visual, presaccadic and cognitive activation of single neurons in monkey lateral intraparietal area. *Journal of Neurophysiology*, 76, 2841–2852.
- Colby, C. L., & Goldberg, M. E. (1999). Space and attention in parietal cortex. *Annual Review of Neuroscience*, 23, 319–349.
- Freedman, D. J., & Assad, J. A. (2006). Experience-dependent representation of visual categories in parietal cortex. *Nature*, 443, 85–88.
- Gottlieb, J., Kusunoki, M., & Goldberg, M. E. (1998). The representation of visual salience in monkey parietal cortex. *Nature*, 391, 481–484.
- Ipata, A. E., Gee, A. L., Goldberg, M. E., & Bisley, J. W. (2006). Activity in the lateral intraparietal area predicts the goal and latency of saccades in a free-viewing visual search task. *The Journal of Neuroscience*, 26, 3656–3661.
- Lewis, J. W., & Van Essen, D. C. (2000a). Mapping of architectonic subdivisions in the macaque monkey, with emphasis on parieto-occipital cortex. *The Journal of Comparative Neurology*, 428, 79–111.
- Lewis, J. W., & Van Essen, D. C. (2000b). Corticocortical connections of visual, sensorimotor, and multimodal processing areas in the parietal lobe of the macaque monkey. *The Journal of Comparative Neurology*, 428, 112–137.
- Moore, T., Armstrong, K. M., & Fallah, M. (2003). Visuomotor origins of covert spatial attention. *Neuron*, 40, 671–683.
- Nakamura, H., Kuroda, T., Wakita, M., Kusunoki, M., Kato, A., Mikami, A., et al. (2001). From three-dimensional space vision to prehensile hand movements: The lateral intraparietal area links the area V3A and the anterior intraparietal area in macaques. *The Journal of Neuroscience*, 21, 8174–8187.
- Oristaglio, J., Schneider, D. M., Balan, P. F., & Gottlieb, J. (2006). Integration of visuospatial and effector information during symbolically cued limb movements in monkey lateral intraparietal area. *The Journal of Neuroscience*, 26, 8310–8319.
- Platt, M. L., & Glimcher, P. W. (1999). Neural correlates of decision variables in parietal cortex. *Nature*, 400, 233–238.

- Reynolds, J. H., & Chelazzi, L. (2004). Attentional modulation of visual processing. *Annual Review of Neuroscience*, 27, 611–647.
- Rizzolatti, G., Riggio, L., Dascola, I., & Umiltà, C. (1987). Reorienting attention across the horizontal and vertical meridians: Evidence in favor of a premotor theory of attention. *Neuropsychologia*, 25, 31–40.
- Roelfsema, P. R., & van Ooyen, A. (2005). Attention-gated reinforcement learning of internal representations for classification. *Neural Computation*, 17, 2176–2214.
- Sugrue, L. P., Corrado, G. S., & Newsome, W. T. (2004). Matching behavior and the representation of value in the parietal cortex. *Science*, 304, 1782–1787.
- Thomas, N. W., & Pare, M. (2007). Temporal processing of saccade targets in parietal cortex area LIP during visual search. *Journal of Neurophysiology*, 97, 942–947.
- Wardak, C., Olivier, E., & Duhamel, J. R. (2002). Saccadic target selection deficits after lateral intraparietal area inactivation in monkeys. *The Journal of Neuroscience*, 22, 9877–9884.
- Wardak, C., Olivier, E., & Duhamel, J. R. (2004). A deficit in covert attention after parietal cortex inactivation in the monkey. *Neuron*, 42, 501–508.