

1 Imaging Expectations and Attentional Modulations in the Human Brain

Maurizio Corbetta and Gordon L. Shulman

1.1 Introduction

Neurobiological research since 1980 has clearly established that visual perception does not depend solely on the neural activity evoked by individual objects, but is also powerfully influenced by contextual sensory information and the behavioral states of the observer. For example, the activity of individual neurons in visual cortex (including area V1) does not depend only on the attributes of the object in the classical receptive field (luminance, contrast, color, orientation, and so on). Rather, the context created by other objects inside or outside the neurons' receptive field (Knierim and Van Essen, 1992; Gilbert, 1996; Ito et al., 1998; Ito et al., this volume; Reynolds and Desimone, this volume) and the level of arousal and interest of the observer are important as well (Wurtz et al., 1980; Motter, 1993). Some of the sensory interactions mediating this context dependence become ingrained during the development of visual cortex, hardwired by the selective pressure of the visual environment over years. Other interactions reflect the ability of visual cortex to learn new visual patterns and configurations over a time scale of minutes or hours (perceptual learning; e.g., Karni et al., 1995; Ito et al., chapter 5 in this volume). In contrast, changes in the behavioral state of the observer can alter the response of visual neurons on the more rapid time scale of milliseconds to seconds. For instance, variations in the level of vigilance occur throughout the day and produce tonic and rather nonselective changes in the level of visual activity (Wurtz et al., 1980; Mountcastle et al., 1987). Cognitive signals such as behavioral goals, expectations, memories, or thoughts, however, produce briefer and more selective modulations of visual activity (for reviews, see Desimone and Duncan, 1995; Maunsell, 1995). This ongoing modulation of visual perception is necessarily selective because at any one time many visual objects compete for awareness, and because multiple cognitive signals can potentially bias perception.

This review focuses on the psychological and neural mechanisms underlying the use of advance visual information, such as knowledge about the color, motion, or location of a target object, for its detection or discrimination in a visual scene. This is a common situation, as when we search for a face or a colored hat in a crowd of people, or when we guess the trajectory of a tennis ball coming our way. In the laboratory this behavior is studied by providing the observer (human or animal) with a cue that carries advance information about some task-relevant target object, and by testing the accuracy or speed of target detection/discrimination in a subsequently presented test display. It is well established that the cue helps the detection/discrimination of relevant objects, and impairs the detection/discrimination of irrelevant objects (Eriksen and Hoffman, 1972; Posner, 1980; Hawkins et al., 1990; Sperling et al., chapter 10 in this volume). Little is known

about how the brain organizes this complex behavior. Where in the brain are cognitive expectations or goals about visual objects or features coded? What is the format of these signals, and how/where in the visual system do they interact with incoming sensory information? What is the effect of such interaction on visual processing, and hence on visual perception?

In the first part of this chapter, we discuss some of the advantages and disadvantages of functional neuroimaging as compared with other neurophysiological methods for measuring neuronal activity during visual tasks. Second, we review some psychological theories as to how cued information influences visual processing. Then, we discuss experiments that demonstrate the effects of focused attention on visual processing. Finally, we present a new functional imaging experiment employing event-related fMRI to dissociate the encoding and maintenance of cue information from its subsequent effect on visual analysis (processing?). Whenever possible, the imaging results will be discussed in relationship to current psychological and neurophysiological findings/ideas on visual selection.

1.2 Tracking Neural Activity during Visual Attention

Tremendous progress has been made in our ability to monitor brain activity at different spatial and temporal resolutions (for reviews, see Wurtz et al., 1984; Hillyard and Picton, 1987; Raichle, 1994). Single-unit recordings in awake, behaving primates, scalp recordings of evoked electrical activity, and imaging of local changes in brain hemodynamics (blood flow or deoxygenation) in human subjects are the main methods used to record activity of the brain during cognitive tasks. Each method provides a view of the brain in action that is biased toward a particular spatiotemporal resolution. Single-unit recordings offer excellent spatial (microns) and temporal (milliseconds) resolution within a cortical region in the monkey, but typically sample only a small fraction of the neurons (cortical areas) involved in a task. Evoked potentials have an excellent temporal resolution, but are biased toward neural activity coming from the surface and have a poor spatial resolution (several centimeters). Finally, functional imaging surveys the whole brain (both cortical and subcortical regions) simultaneously with a spatial resolution of a few millimeters, but with a temporal resolution that is coarser (several seconds) than real-time neural activity. An additional problem with imaging is that the local hemodynamic signals which are recorded represent only an indirect measure of neural activity, and the details of the coupling between blood vessels and neurons are presently unknown. There is good evidence, however, that hemodynamic signals precisely colocalize (within hundreds of microns) with neuronal activity recorded by single units in monkey's visual cortex (Ts'o et al., 1990). Furthermore, hemodynamic signals vary linearly within certain parameter ranges with several psychophysical visual functions (e.g., contrast sensitivity; Boynton,

1996) and cognitive variables (e.g., load during working memory tasks; Nystrom et al., 1998).

More direct information about the coupling of neuronal activity and hemodynamics will become available once fMRI for monkeys is more fully developed (Stefanacci et al., 1998; Logothetis et al., 1999). An important recent development in human fMRI is the capability to record focal changes in blood oxygenation caused by single sensory, cognitive, or motor events (event-related fMRI; Zarahn et al., 1997; Dale and Buckner, 1997; Friston et al., 1998; Rosen et al., 1998). This constitutes a significant improvement over earlier PET or fMRI designs in which activity had to be recorded, averaged, and displayed over many trials (blocked design). Event-related fMRI allows the randomized presentation and analysis of separate trial types. Recently, our laboratory has further extended this approach to the analysis of different events within a trial (Ollinger et al., 1998). For example, regions of activation related to the presentation of a cue can be now differentiated from regions related to the presentation of a test stimulus or a motor response (Shulman et al., 1999). This method is very helpful for tracking the slow temporal evolution of the hemodynamic signals.

1.3 Psychological Theories about the Selection of Simple Visual Features

Psychological studies have characterized the experimental conditions under which advance information about simple visual attributes facilitates the perception of relevant stimuli, and impairs the perception of irrelevant stimuli (Pashler, 1998). Highly discriminable cues such as location, color, and size are very effective in facilitating perception of subsequently presented stimuli. The facilitation is greatest when the discrimination is difficult or in the presence of irrelevant stimuli. Under optimal selection conditions there is little evidence of processing of irrelevant stimuli (but see Braun et al., chapter 11 in this volume). These findings are consistent with the notion that cue information interacts with processes involved in the sensory and/or decisional analysis of both relevant and irrelevant stimuli, and that this interaction relatively facilitates the processing of relevant stimuli. It is currently debated whether an object can be selected directly on the basis of information about its intrinsic features such as color, size, or direction of motion, independently of its location, or whether location analysis is indispensable for object selection (Moore and Egeth, 1998; Shih and Sperling, 1996).

Pashler (1998) reviews two general accounts of the effects of cues on the analysis of incoming sensory information. On one account, advance information enhances the sensory processing of the signal (signal enhancement mechanism). On the other account, advance information allows observers to disregard channels containing noise, that is, information that could degrade a perceptual decision if allowed to influence the decision. The ability

to disregard information from irrelevant channels can improve perception for merely statistical reasons by minimizing the chance of a false alarm, that is, erroneous detection based on noise information. Noise suppression might be implemented through a variety of mechanisms: raising the threshold for decision and/or suppressing sensory information in the irrelevant channel. A related issue, which has driven much of the psychological research on attention since 1960, is at what level of processing selection mechanisms operate (Deutsch and Deutsch, 1963; Duncan, 1980; Broadbent, 1982; Treisman, 1969). Selection mechanisms (both signal enhancement and noise suppression) might work at early (anatomically and temporally) or late stages of visual processing, that is, respectively influence sensory and/or decisional stages of analysis.

Can we use these ideas to generate predictions about patterns of neural activation related to the encoding, maintenance, and use of visual expectations, and their influence on sensory/decisional visual processes? In terms of imaging research, a straightforward prediction about putative mechanisms of visual selection is that the site of neural modulation should differ, depending on the underlying mechanism. In particular, a mechanism that enhances signal should involve predominantly task-relevant pathways, whereas a mechanism that suppresses noise should at least partly involve irrelevant pathways. For example, if a subject is required to attend to discriminate a feature of a stimulus (e.g., its speed) while ignoring other features (e.g., its color or shape), most of the signal should be present in motion-sensitive regions, whereas most of the noise should come from color- and shape-sensitive regions. Modulations in motion-sensitive areas would support signal enhancement accounts; modulations in color- and shape-sensitive areas would favor noise suppression accounts (figure 1.1). Similarly, if a subject is required to attend to a location in the visual field and to discriminate the orientation of stimuli presented at attended or unattended locations, modulations at either attended or unattended parts of retinotopically organized cortical areas would support, respectively, signal enhancement or noise suppression accounts. Finally, modulations that occur early in the visual system (e.g., area V1) are unlikely to alter decisional stages of analysis. In this case it is important to show that modulations occurring anatomically early in the visual system do not reflect temporally late feedback signals from postperceptual levels.

An important caveat in considering a neural implementation of psychological mechanisms is that the visual system comprises more than thirty visual cortical areas that are hierarchically organized and reciprocally linked by several hundred connections (for a review, see Van Essen and DeYoe, 1995). Hence, psychological distinctions (and models) about information-processing stages—such as input, sensory analysis, stimulus identification, and decision—must take into account the great complexity of the underlying neural implementation. In the brain, different stages of processing are likely to overlap across multiple areas; decisions may be generated at different levels, depending on task demands;

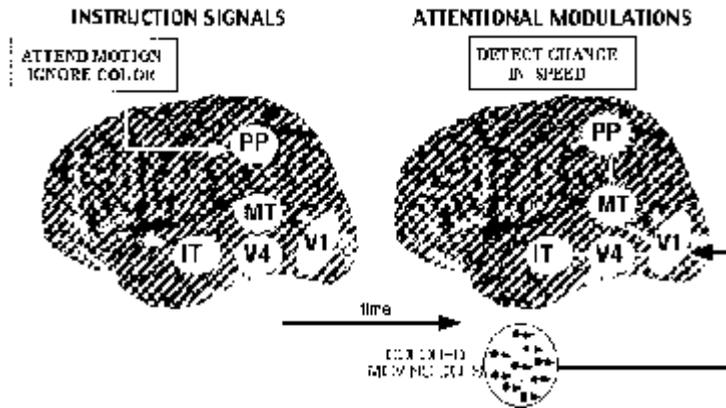


Figure 1.1

Expectation signals and attentional modulations. Expectation signals are neural signals that encode/maintain visual expectations about a visual attribute (perceptual set) (*left*). PP, posterior parietal cortex; MT, middle temporal cortex; V4, area V4; IT, inferior temporal cortex; V1, area V1. The preactivation by expectations of relevant pathways (solid line to PP) is consistent with capacity allocation (signal enhancement) mechanisms; the preactivation of irrelevant pathways (broken line to IT) is consistent with noise suppression mechanisms. Expectation signals interact with incoming visual information and produce attentional modulations. The neuro-imaging evidence during attention to stimulus features suggests modulation of task-relevant pathways (e.g., MT and PP).

and the flow of information is likely to be bidirectional from lower to higher levels, and vice versa. Moreover, attentional effects on perception and behavior may well reflect modulations at multiple cortical levels from sensory to motor centers.

1.4 Expectation Signals and Attentional Modulations

Focusing attention on simple visual features not only facilitates the perception of relevant stimuli in a visual scene, but also powerfully modulates the neural activity in visual cortex evoked by those stimuli. Logically, it seems important to separate processes (or neural signals) relating to the establishment and maintenance of visual expectations established by a cue (which we call *expectation signals*; “template signals” in Desimone and Duncan, 1995), on the one hand, from processes/signals reflecting selective modulation of sensory activity evoked by a test stimulus (which we call *attentional modulations*), on the other hand. These two kinds of signals presumably have different time courses, since expectation signals must precede attentional modulations. They may also have different spatial distributions within the brain. For example, whereas attentional modulations may emphasize visual areas involved in the analysis of the test stimulus, expectation signals may involve

both nonvisual areas, necessary for their encoding and maintenance in time, and visual areas, which receive (and possibly maintain) them (figure 1.1).

Finally, a crucial aspect of human cognition is that similar behavioral goals and expectations can be generated from information that occurs in widely different formats. For instance, the search for a red car in a parking lot might be initiated verbally (e.g., a friend's comment, "My car is the red one"), by information stored in long-term memory (e.g., I know that my car is red), or by recently presented sensory information (e.g., "Where is the red car we saw enter the parking lot?"). In all these cases, an expectation is established that guides the search for the red car. Hence, similar expectation signals may be expected for cues presented in different sensory modalities (e.g., vision, audition) or formats (e.g., iconic, symbolic, linguistic).

Although many studies have demonstrated the existence of attentional modulations in both monkey and human visual cortex (for reviews of primate studies, see Desimone and Duncan, 1995; Maunsell, 1995; and Motter, 1993, 1994a, 1994b; for human studies, see Hillyard and Picton, 1987; Corbetta et al., 1991; Dupont et al., 1993; Haxby et al., 1994; Mangun et al., 1993; Beauchamp et al., 1997; Buckner et al., 1997; Tootell et al., 1998; Wojciulik et al., 1998; Martinez et al., 1999; Shulman et al., 1999; Gandhi et al., 1999), little is known about how cognitive expectations and goals are encoded and maintained in the brain. Below, we first review what has been learned through neuroimaging about attentional modulations in the visual system, and then consider a new experiment aimed at isolating expectation signals.

1.4.1 Attentional Modulations

Imaging studies since 1990 have clearly demonstrated several general rules about attentional modulations in the visual system.

First, directing attention to simple visual features such as color, motion, shape, and location, or more complex objects such as faces, words, or buildings, modulates activity in task-relevant pathways, that is, in visual areas specialized for processing the selected visual attribute or object. For example, in our original positron emission tomography (PET) study (Corbetta et al., 1990, 1991) subjects performed a match-to-sample task on a random display of colored moving bars. In different scans, subjects either attended to/discriminated a particular kind of change of the stimuli (selective attention to either color, motion, or shape), or attended to/discriminated any kind of change (divided attention between color, motion, and shape), or simply detected stimulus onset (passive viewing). Discrimination thresholds were lower with selective than with divided attention, confirming that selective attention enhances perception. PET measurements demonstrated higher activity in the visual system during selective attention as compared with passive viewing or divided attention. In other words, activity was higher when attentional re-

sources were concentrated on one particular pathway than when they were presumably spread across multiple pathways. For example, during motion discrimination, activity was enhanced in extrastriate visual areas known to be sensitive to visual motion (namely, in a collection of areas called MT+ that includes area MT/V5).

The selective modulation of task-relevant pathways while focusing attention on a particular stimulus feature has been confirmed by other studies. O'Craven et al. (1995) found a 27% increase in blood oxygenation level-dependent (BOLD) signal in area MT+ when subjects attended in the same display to moving versus static random dots. Beauchamp and colleagues (1997) found that the BOLD signal dropped by 56% in MT+ and the intraparietal region (another motion-sensitive region) when subjects diverted attention from stimulus speed to stimulus color in a colored random dot display. The signal fell by 150% when subjects diverted attention from the random dot display to a small central fixation point in order to detect changes in its luminance. This study therefore indicates that activity in the same visual area can be modulated by attending to the task-relevant feature and/or location, and that the two effects may be additive.

In addition, many other studies have found an effect of cueing location alone (Corbetta et al., 1993; Vandenberghe et al., 1996, 1997; Woldorff et al., 1997; Tootell et al., 1998; Brefczynski and DeYoe 1999). Directing attention to the location of flashed stimuli to be detected/discriminated modulates activity in visual areas that code for stimuli at the attended location. Finally, other studies have found modulations in task-relevant pathways by directing attention to more complex stimuli such as faces or buildings (Haxby et al., 1994; Wojciulik et al., 1998). For example, Haxby and colleagues, using PET, observed more ventral occipital activity during face processing and more dorsal parietal activity during location discrimination on the same set of visual stimuli. Wojciulik and colleagues found that activity in a ventral region specialized for face perception was 145% stronger when subjects processed faces than when they processed objects on the same set of stimuli.

It is less clear if the described attentional modulations reflect a relative increase of activity in task-relevant pathways or a relative suppression in task-irrelevant pathways. This is important in deciding between signal enhancement vs. noise suppression accounts. Direct comparisons between attended versus unattended conditions are not helpful because the corresponding modulation may reflect either a relative increase or a relative decrease in activity. Instead, it is necessary to gauge the sign of the modulation with an independent control condition. The approach used in many experiments has been to introduce a passive viewing condition, in which subjects are presented with the same set of stimuli and have to respond with a key-press to the onset of each display. The word "passive" is unfortunate because it may suggest large differences in the level of vigilance from the "active" tasks. Instead, a passive viewing control may be quite effortful, because subjects have to

maintain accurate eye fixation, detect the onset of temporally asynchronous visual stimuli, and prepare/execute an appropriate detection response. A more appropriate term may be “detection” task. The logic behind a passive viewing task is to provide an independent sensory-motor baseline in which selective attention is not explicitly and endogenously directed toward certain stimuli or features. Admittedly, the control is problematic because attention is unchecked. On the other hand, it is not unreasonable to think that in this task attention is mostly drawn by the physical properties of the display, and thus evenly distributed across stimuli or stimulus features.

In our original experiment (Corbetta et al., 1990, 1991), we found no modulation in regions coding for the irrelevant feature as compared with a passive viewing control or a selective attention condition. In other experiments, which involve focusing attention toward one visual field location and bilateral presentation of target or distractor stimuli, attentional modulations have been found contralateral to the attended side (the side representing the attended location). In contrast, little or no modulation has been observed ipsilateral to the attended side, contrary to what one would expect if the role of spatial attention was to filter out irrelevant information from unattended locations (e.g., Vandenberghe et al., 1996, 1997; Woldorff et al., 1997). These effects have been correlated with improved psychophysical performance and early temporal modulations on scalp electrical potentials (for a review, see Hillyard et al., 1998). Recent elegant experiments by Tootell and colleagues (1998) and Brefczynski and DeYoe (1999) found that directing attention toward multiple locations in the visual field enhances cortical activity at the corresponding locations of retinotopically organized visual areas, including area V1 (see also Gandhi et al., 1999; Martinez et al., 1999; Somers et al., 1999).

These findings show directly that attending to a particular visual location enhances visual performance at that location by modulating activity at the corresponding points of retinotopically organized visual cortex. Tootell and colleagues (1998) also reported that the allocation of attention to peripheral field locations produced a suppression of activity in adjacent cortical regions representing the fovea, which was unattended during the experiment. Interestingly, the region of cortical suppression did not extend to other cortical locations that were coding for distractor (irrelevant) stimuli presented in other visual quadrants, either in the same or in the opposite hemisphere. Hence, the suppressive modulation is spatially incongruent with the source of noise, and it is unclear whether this reflects the neural correlate of a noise suppression mechanism.

Another common finding is that attentional modulations involve multiple visual areas within a task-relevant pathway, including at times primary visual cortex (area V1). In our original experiment, we found modulations at early (near calcarine sulcus, primary visual cortex), intermediate (MT+), and late (parietal) stages of the visual hierarchy. Similarly, Buechel and colleagues (1998) found modulations in areas V1/V2, V3, V5, and parietal

cortex during attention to motion as compared with passive viewing. Haxby and colleagues (1994) showed that attention to faces modulated a large swath of tissue in ventral occipital cortex, overlapping with the multiple foci identified by Corbetta and colleagues (1991) for shape discrimination. These are likely to correspond to intermediate and higher levels of the ventral stream of processing (areas V4, TE, TEO). Recently, there have been elegant demonstrations that attention to specific visual field locations enhances activity in corresponding retinotopic locations of multiple early-intermediate visual regions (Kastner et al., 1998; Tootell et al., 1998; Brefczynski and DeYoe 1999), including area V1 (Tootell et al., 1998; Somers et al., 1999; Heeger et al., chapter 2 in this volume).

These imaging findings indicate that attentional modulations occur within a sensory pathway at early, intermediate, and late stages of processing. It is therefore likely that both perceptual and decisional processes are modulated. The predominance of modulation in task-relevant pathways suggests that attention may influence the signal more than suppressing the noise from irrelevant channels. The modulations might reflect top-down expectation signals that act directly at multiple levels of a sensory pathway (Olshausen et al., 1993). Alternatively, the occurrence of modulations at different levels may reflect a cascade effect due to existing connectivity caused by a top-down signal that primarily acts at a single level in the hierarchy of visual areas (Tsotsos et al., chapter 14 in this volume).

Open questions for future research include the quantification of these attentional modulations (in terms of both volume and magnitude) across multiple visual areas, their relationship to individual psychophysical performance, and how they relate to the attentional modulations recorded with other methods. For example, Treue and Maunsell (1996) reported a neuronal enhancement for attended stimuli of 86% and 113% in areas MT and MST, respectively, during a speed discrimination task in which attended and unattended targets were placed within the same receptive field. The size of this modulation is comparable with the one obtained in imaging experiments on attention to motion (e.g., 56–150% in Beauchamp et al., 1997).

However, there are also some discrepancies between methods. Heeger and colleagues (this volume) report strong attentional modulation in area V1 with fMRI during spatial attention. For many years, modulations of area V1 have been difficult to demonstrate by single-unit recordings in monkey. Similar attentional modulations in area V1 have been reported in fMRI studies from other laboratories (Martinez et al., 1999; Somers et al., 1999), and represent a robust and easily replicable finding. Martinez and colleagues compared attentional modulations in area V1 observed with different methods (fMRI and ERP), and suggested that the hemodynamic modulations may reflect temporally late feedback signals into area V1 from higher visual areas. They reached this conclusion because no attentional modulation was evident on the earliest visually evoked electrical potentials.

It is possible, therefore, that single-unit recording and functional imaging may sample neuronal signals in area V1 on different temporal scales.

Another puzzle is that attentional modulations in single-unit experiments are much stronger when target and distractors are placed within the same receptive field. For example, the neuronal enhancement reported by Treue and Maunsell (1996) fell from 86% to 19% in area MT, and from 113% to 40% in MST, when the distractor stimulus was placed in the opposite visual hemifield rather than next to the target stimulus in the receptive field of the recorded neuron (Maunsell and McAdams, this volume). The weakness of attentional modulations when stimuli are not positioned within the same receptive field is also a common finding in the ventral system (including area IT, where receptive fields are large and span both visual hemifields) (Moran and Desimone, 1985). In contrast, very strong attentional modulations can be obtained with both evoked potential and imaging methods by placing target and distractor items in opposite hemifields. The argument usually put forward in the single-unit literature is that modulations at the single-neuron level are best observed when there is competition between stimuli in the same receptive field (e.g., Desimone and Duncan, 1995).

This raises the paradox that correlates of neural competition between stimuli far apart on opposite sides of the vertical meridian would be observed only in areas with bilateral receptive fields. It is unclear whether the brain has such neurons, but there is strong psychophysical, imaging, and evoked potential evidence indicating that spatial competition between stimuli is resolved rather early in the visual system (possibly as early as area V1 under certain conditions). Therefore significant differences must exist between modulations observed at the population (e.g., imaging and evoked potential) and the single-unit level. Although there are many potential explanations for this discrepancy, one possibility is that the extracellular recording of the spiking activity of single neurons in visual cortex does not capture some critical neural codes reflecting the allocation of attention to spatial locations. For example, typical single-unit recordings would not detect systematic temporal correlations in the firing of neuronal populations within an area, which might underlie important behavioral, surface electrical, or hemodynamic modulations related to attention (Engel et al., 1991). It will be critical in the near future to develop more quantitative models of the relationship between vascular and neuronal signals. This will require comparing in the same animal, and under the same task conditions, hemodynamic signals, single-unit firing rates, and other measurements characterizing larger neuronal populations (e.g., multiunit recording, cross-correlation analysis).

1.4.2 Expectation Signals

An important ambiguity in all imaging experiments reviewed so far is their failure to distinguish between visual expectation and the effect of that expectation on visual pro-

cessing. Expectations (e.g., what type of stimuli are task-relevant) are typically induced by top-down signals that are triggered by an appropriate cue, and that encode and maintain the relevant perceptual set. These neural signals must be separate from the modulations they induce on the sensory activity evoked by a subsequent stimulus (attentional modulations). Blocked PET and fMRI designs average activity over many trials, and therefore blur differences between processes that are active at different times within a trial. The ability to distinguish between earlier expectation signals and later attentional modulations is critical for the question of how visual expectations are implemented in the brain, and whether they resemble previously described mechanisms of signal enhancement or noise suppression. Signal enhancement models would predict expectation signals to be present predominantly on task-relevant pathways, whereas noise suppression models would predict expectation signals to be localized predominantly on task-irrelevant pathways. The reviewed imaging evidence, which finds predominant modulations in task-relevant pathways, is ambiguous in this respect because expectation signals, visual responses, and related attentional modulations were temporally confounded.

At present, little is known about the neural basis of visual expectations (or expectation signals) even from a single-unit perspective. Several neuronal correlates of expectation signals have been described. Many authors have reported an increase in the baseline firing rates of neurons in several brain regions when the monkey can anticipate the location of an upcoming target stimulus (Bruce and Goldberg, 1985; Colby et al., 1996; Luck et al., 1997). The effect is endogenous and spatially selective, that is, the attended location coincides precisely with the neuron's receptive field. Luck and colleagues (Luck et al., 1997; Reynolds and Desimone, chapter 7 in this volume) have distinguished the tonic modulation of the baseline response from a time-locked neuronal enhancement of the visual response triggered by the presentation of a target stimulus. They have argued that the tonic baseline increase reflects a selective biasing signal for location (expectation signal), whereas the enhancement is the result of a change in the gain of the sensory response caused by the biasing signal (the attentional modulation of the test stimulus).

Other studies have recorded cue-specific neural activity from several extrastriate and prefrontal areas during the cue period of delayed match-to-sample tasks. Typically, the animal is shown a sample object (or cue, which generates the expectation signal) and has to determine whether it matches subsequent objects (Fuster, 1973; Chelazzi et al., 1993; Miller et al., 1996). Interestingly, cue-related activity can be recorded in visual cortex (e.g., in area V4) even when the cue information is delivered in a tactile format (Haenny et al., 1988). Such delay activity therefore represents a neuronal correlate of visual expectations about upcoming stimuli. However, in the presence of intervening distractor stimuli (i.e., stimuli that do not match the cue), cue-specific delay activity is disrupted in extrastriate visual cortex but is maintained in prefrontal areas (Funahashi et al., 1993; Miller et

al., 1996). Desimone and Duncan (1995) have therefore proposed that these prefrontal areas are the source of expectation signals that produce attentional modulations in posterior visual areas.

Overall, these studies clearly demonstrate that expectation signals generated by the presentation of visual cues are present in visual cortex prior to the presentation of a test stimulus, which is consistent with psychological theories in which cue information preactivates sensory channels. However, these signals have been recorded from one visual cortical region at a time (but see Miller et al., 1996), so that expectation signals and attentional modulations have not been characterized simultaneously over the entire brain on the same task. Furthermore, one cannot rule out the possibility that expectation signals generated through iconic cues (as in most match-to-sample paradigms) might be partially confounded with visual activity. Finally, we can generate expectations about visual attributes or scene in a variety of ways: through vision, audition, or touch of a sample attribute/object, verbal information derived either from outside cues (e.g., visual or auditory words) or internal processing, verbal or visual memory, imagery, and so on. How are all these different sources of expectation/bias organized and funneled in the visual system, in formats that are compatible with those of cortical visual areas? Is a common set of brain systems (e.g., working memory systems) involved in maintaining on-line different expectations?

1.4.3 Dissociating Expectation Signals and Attentional Modulations with Event-Related fMRI

The main purpose of this experiment (fully reported in Shulman et al., 1999) was to develop a method to separate BOLD signals related to the encoding/maintenance of expectation signals, from their effect (attentional modulation) on a subsequent visual response evoked by a test stimulus. Therefore, we used psychophysical tasks in which expectation signals could be generated in one condition, but not in a control condition, while matching basic sensory-motor variables. To avoid concerns about sensory activations produced by iconic cues, we elected to instruct subjects through symbolic cues, which require active transformation and do not contain motion energy. The task is a cued motion coherence detection task, modified after Ball and Sekuler (1981) and Newsome and colleagues (Britten et al., 1993). Ball and Sekuler have previously shown that stationary direction cues improve motion detection. Newsome and colleagues have extensively used a motion detection coherence task to understand the contribution of area MT in monkey to motion processing. The motion system is the best understood sensory system in both monkeys and humans. Several motion-selective regions have been recently described in the human brain (for review, see Tootell et al., 1996). This information allows us to label more precisely areas in which attentional modulations are found, and thus better define the putative stage of attentional selection.

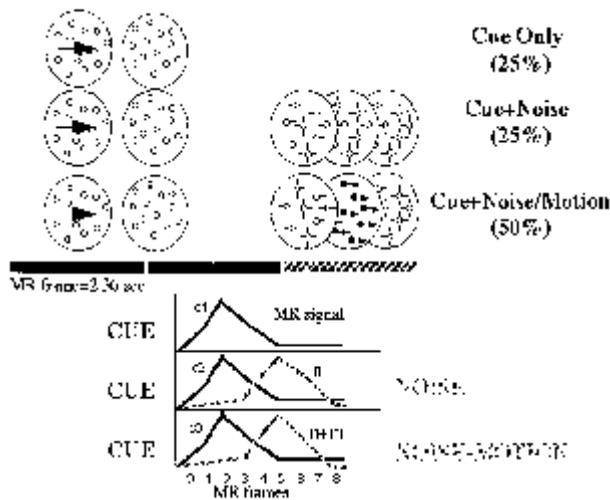


Figure 1.2

(*Top*) Stimulus sequence of a cue trial, a cue + noise trial, and a cue + noise/motion trial. The circular aperture was centered on the fovea and was 3.25° in diameter. In the actual display, no solid line defined the outer limit of the aperture. The cue was presented for 1600 ms. Dot density ($n = 50$) and size were the same for the cue and test periods. Dynamic noise is indicated by the starred dots, and coherent motion is indicated by the arrowed dots (speed 4.3° per second). The motion could be in any of eight directions (given by 45° increments from upward motion). 50% of the trials during the MR session were cue + noise/motion trials, 25% were cue + noise trials, and 25% were cue trials, with the three trial types randomly mixed. The temporal relationship of each trial event to the corresponding MR frame (scan) is indicated. The MR frame (2.36 s) is the time necessary to take a snapshot of the whole brain. (*Bottom*) Linear decomposition of the hemodynamic response function. The MR signal during a cue + noise/motion trial is modeled as composed of two functions, one for the cue period (c^3) and one for the noise plus motion ($n + m$) period. Direct parameter estimation is performed with a linear model for these two functions at each MR frame. Cue trials provide a direct estimation of the MR signal during the cue period in isolation (c^1). Cue + noise trials provide a direct estimation of the MR signal during the cue (c^2) and noise (n) periods in isolation. This information is used to estimate components during the cue period across all trial types (c^{123}), the noise (n), and the noise plus motion ($n + m$) periods.

In the fMRI experiments we ran three types of scans: *directional cue*, *passive cue*, and *motion localizer*. During *directional cue* scans three types of trials were randomly intermixed (figure 1.2). All trials began with a stationary arrow cue, presented at fixation on a random field of static dots for 1600 ms. The arrow cued one of eight possible directions of subsequent motion. Following the extinction of the cue, subjects held fixation on a small cross at the center of the random field of static dots for 3120 ms, for a total cue period of 4720 ms. Therefore, most of the cue period was spent encoding and holding in mind the direction of the arrow. The first type of trial ended following the completion of the cue period (*cue trials*, 25% of total). These were necessary to isolate expectation signals during the cue period. In the second and third types of trials, the static dots were

randomly replotted following the cue period (once each display frame), producing dynamic random noise during the test period. In *cue + noise/motion trials* (50% of total), a percentage of the dots moved coherently (in the direction of the earlier arrow cue) during a brief interval (300 ms) at some randomly chosen time within the test period. Subjects detected the motion by pressing an MR compatible key-press with their right hand. The percentage of coherent motion was adjusted in each subject to yield just above threshold performance (d' approximately 2.0). In *cue + noise trials* (25% of total) only dynamic noise was presented (i.e., dots never moved coherently), and subjects had to withhold a response. Each trial lasted on average 4720 ms, and the intertrial interval was randomly varied between 4720 ms, 7080 ms, and 9440 ms.

During *passive cue* scans, subjects passively viewed trial sequences that were identical except that the cue was a filled square (rather than an arrow) and that subjects did not make a response. *Passive cue* scans served as a control condition in which the cue generated no expectation about direction of motion. These scans also provided a control for the sensory activity evoked by the onset of the cue and the test stimulus. Differences in magnitude of the response to the test stimulus during directional and passive cue trials provided a measure of (attentional) modulations produced by a motion set generated through directional arrow cues during the detection of coherent motion.

Finally, to determine which areas were activated by sensory motion, subjects received *motion localizer* scans in which periods of continuous radial dot motion were alternated with control periods in which the dots were stationary. Subjects looked at the display and maintained fixation on a central crosshair. Areas active during radial motion were used for comparison with the areas that carried expectation signals and attentional modulations.

To verify that subjects were using the directional cue, we conducted a separate psychophysical session in which directional cue trials were randomly intermixed with neutral cue trials (in which a plus sign cued all eight possible directions of motion). Psychophysical models of motion selection suggest that a neutral cue does not instruct any particular direction but a general set for motion. Behaviorally, subjects were faster and more accurate when they used directional cues than neutral cues, indicating that the arrow cue was successful in generating an expectation for direction of motion that facilitated the perception of subsequently presented coherent motion.

The BOLD responses initiated during the cue and test periods were estimated with a linear regression model assuming that the MR signal on any frame is the linear sum of different components. No assumptions were made about the shape of the hemodynamic response function. Previous work using event-related fMRI design and rapid presentation rate has successfully estimated the BOLD response to different *trials* by randomly jittering the intertrial interval (e.g., Dale and Buckner, 1997). To estimate the response to different events *within a trial* (e.g., BOLD responses to the cue and to the test stimulus), it is also necessary to present a small percentage of “catch” trials in which only the first component

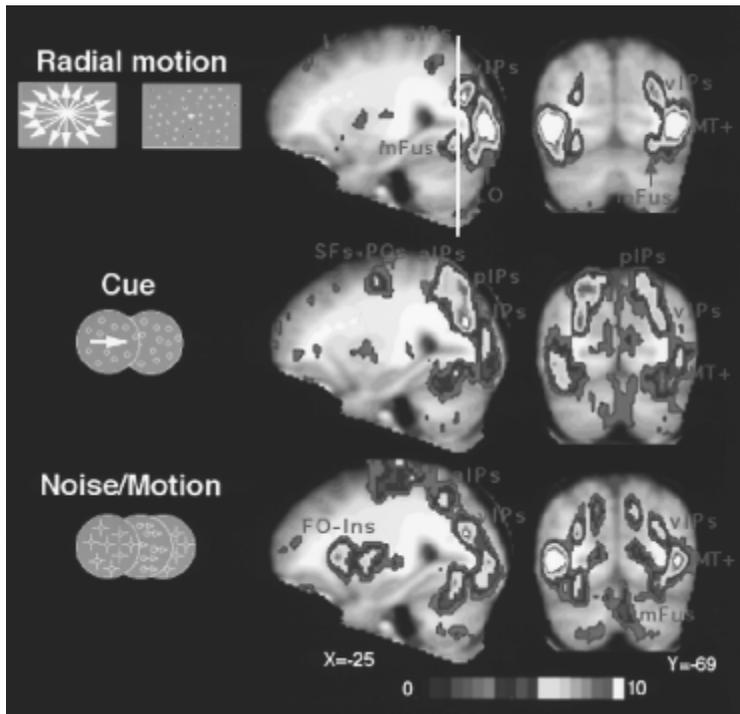


Figure 1.3

Group z-maps for activations during radial motion (*top row*), the cue period of trials involving a directional cue (*middle row*), and the noise/motion period of trials involving a directional cue (*bottom row*). A sagittal slice of the left hemisphere is shown in the left column, a coronal slice in the right column. The white line through the sagittal slice in the top left panel shows the location of the coronal slice. The color scale represents the z-score of the activation, and all displayed pixels have passed a multiple comparison procedure that includes a Bonferroni correction for the number of hemodynamic response functions used to generate the z-map. aIPs, anterior intraparietal sulcus; pIPs, posterior intraparietal sulcus; vIPs, ventral intraparietal sulcus; SFs-PCs, superior frontal-precentral sulcus; LO, lateral occipital; FO-Ins, frontal operculum-insula; mFus, mid-fusiform gyrus. (See plate 1 for color version.)

is present (cue-only trials). A separate validation of this new method has been presented by Ollinger and colleagues (1998).

The top row of figure 1.3 (see also plate 1) shows areas active for radial sensory motion. They include the mid-fusiform gyrus (mFus), lateral occipital, middle temporal area (MT+), ventral and anterior intraparietal regions (vIPs, aIPs), and cortex near the precentral sulcus (PrCs).

Some of the same motion-sensitive regions were active during the cue period following the presentation of a (static) directional cue (figure 1.3, middle row). They included the anterior and ventral IPs region, the ventral portion of area MT+, lateral occipital, and the

precentral region (at the intersection with the superior frontal sulcus on the left). This activity did not simply reflect the sensory response evoked by the onset of the cue stimulus, because it was significantly stronger during directional than passive cues. Other non-motion-sensitive regions were also active during the cue period. A region in the posterior portion of the IPs was uniquely active during the cue period. A region in the anterior fusiform gyrus (antFus), located more laterally, ventrally, and anteriorly to the one active for radial motion, strongly responded during the cue period for both directional and passive cues. The location of this region is near cortex previously found active during face and shape processing (Haxby et al., 1994; Wojciulik et al., 1998), and may be involved in the initial shape analysis of the static cue. This interpretation implies that some shape analysis is conducted on both directional and passive cue trials, that is, both when the cue is relevant and when it is irrelevant.

Figure 1.4 shows averaged (across subjects) time courses in motion-sensitive (aIPs and MT+) and non-motion-sensitive regions (pIPs, antFus) during the cue and test periods for directional and passive trials. Note that anterior IPs and MT+ show a separate response during both the cue and the noise/motion (or noise, not shown) periods, whereas posIPs is active only during the cue period. The magnitude of the response is significantly stronger during directional than during passive trials, indicating that these signals do not simply reflect the sensory onset of the cue. These time courses clearly demonstrate that this method can separate processes within a trial. They also provide information about the relative role of each region in the processing of the cue. For instance, activity in parietal cortex (anterior and posterior IPs) was more sustained than in MT+ or in other occipital regions (e.g., antFus). Also, activity in antFus was the strongest across all visual areas during directional and passive cues.

One possible model is that the static shape cue was initially encoded in the antFus region (active in both cue conditions), and transformed into a suitable motion signal via early motion regions (e.g., MT+) that showed an early response during the cue period. The expectation signal for motion was then maintained in regions of the intraparietal sulcus during the cue period. These parietal regions could also participate in the encoding phase. This model suggests that the intraparietal regions are the source of top-down signals which encode expectations about upcoming visual motion in this task. This result is evocative of recent results of Maunsell and colleagues, who found that neuronal modulations for motion discrimination within the intraparietal sulcus are strongly dependent on the information carried by the cue (see Maunsell and McAdams, chapter 6 in this volume). Surprisingly, no prefrontal activity was localized during the cue period (see below), as one might expect on the basis of reports of cue-related delay activity in prefrontal cortex during match-to-sample tasks (Funahashi et al., 1993; Miller et al., 1996), although a left precentral region (SFs-PrCs) showed cue-related activity. Therefore, any link between

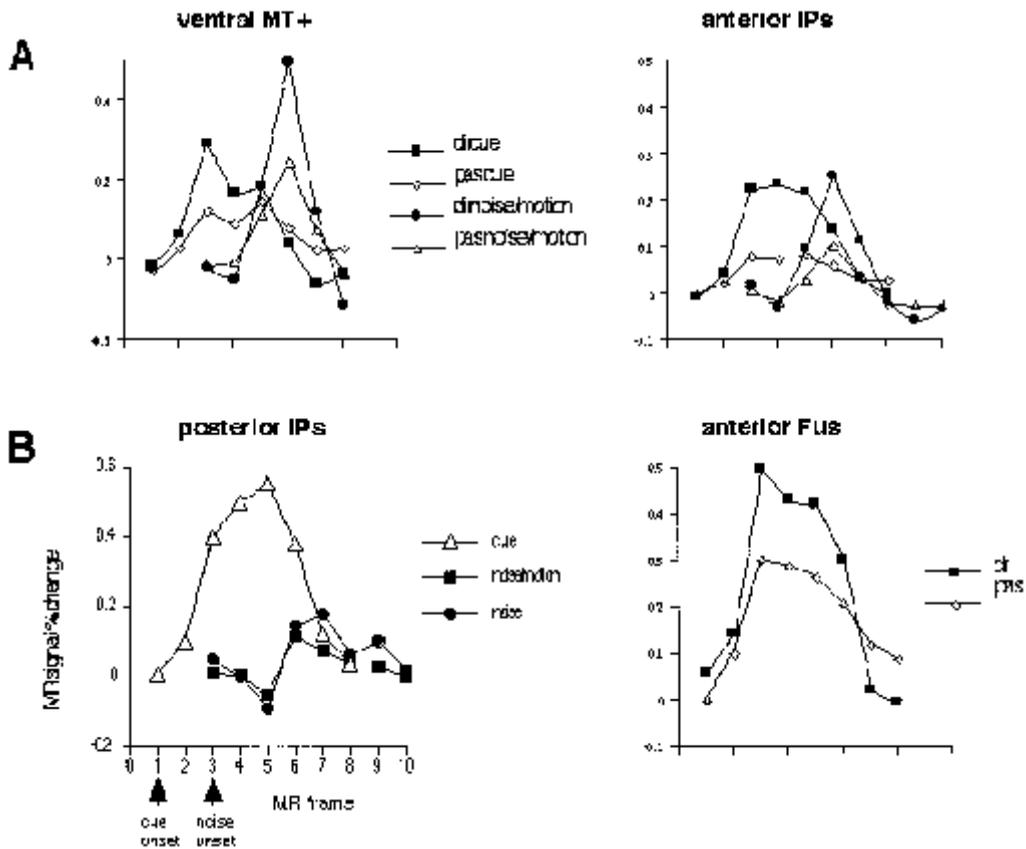


Figure 1.4

(A) Group time course of activations in two motion-sensitive regions (left vMT+, left antIPs) during the cue and test periods for directional and passive cues. Time courses are averaged over 3D regions of interest formed from voxels that were significantly activated during both periods. Both regions show a response during the cue period (after taking into account the hemodynamic delay), followed by a second response during the test period. Note difference in response magnitude between directional and passive trials, and sustained cue-related activity in antIPs but not vMT+. (B) Group time course of activation in non-motion-sensitive regions left posIPs during the cue and test periods, and right antFus during cue period only. posIPs was significantly active only during the cue period, and showed a sustained response. antFus showed strong responses both during directional and passive cues. Compare passive response in antFus with those recorded in vMT+ and antIPs. The antFus may be critical for the analysis of the shape of the cue during both trial types. AntIPs, anterior intraparietal sulcus; posIPs, posterior intraparietal sulcus; antFus, anterior fusiform gyrus; vMT+, ventral area MT complex.

posterior parietal and anterior regions during the cue period might have involved this precentral region. In any case, activity was still significantly more sustained in pIPs than in SFs-PrCs.

All motion-sensitive areas were activated at the presentation of the test stimulus (figure 1.3, bottom row, and time course in figure 1.4), and showed a stronger response to the test stimulus during directional than passive trials. This attentional modulation partly reflects the effect of the directional cue on the sensory processing of the test stimulus. Many other regions, more anterior in the brain, were also activated during the test period. Some regions, such as primary sensory-motor cortex (and the parietal operculum, or area SII), were related to the execution of the key-press response, and showed higher activity during response (hit and false alarm) than no response (misses and correct rejections) trials. Other regions, such as SMA, anterior cingulate, and basal ganglia, showed activations that began after the onset of the noise display, and were present for all trial types (hits, false alarms, misses, correct rejections). Activity in these regions may reflect preparatory motor signals initiated by the onset of the noise display, followed within a few hundred milliseconds by a target triggering a key-press response or by processes involved in searching a target in a noisy display.

A similar pattern was observed in dorsolateral prefrontal cortex (DLPFC) and frontal opercular regions. These regions have been activated during working-memory experiments, and are considered, on the basis of neuronal recordings, to be the source of biasing signals to the visual system. It is therefore important to consider the time course of activation in this region during cue and test periods. Figure 1.5 shows time courses in right and left DLPFC. The location of this region is analogous to the one reported in many working-memory experiments, and corresponds to Brodmann area 46. Note that DLPFC is silent during the cue period but strongly responds during the test period. The type of decision or execution of a response does not affect the magnitude of the response. These findings demonstrate that DLPFC is not universally involved in storing visual expectations which guide visual behavior, and that under certain circumstances this information can be stored in posterior visual regions. The absence of DLPFC activity during the cue period is consistent with a view in which prefrontal cortex is engaged only when stored information is actively manipulated or transformed. An alternative possibility is that DLPFC might be driven by more complex storage operations which perhaps involve a longer delay, linguistic information, or a more precise representation of the cue (e.g., during a match-to-sample task).

In conclusion, this experiment demonstrates that expectations about upcoming visual motion generated by a static symbolic cue preactivate motion pathways involved in the analysis of subsequently presented motion stimuli, as well as non-motion-sensitive regions. The intraparietal sulcus regions are the source of top-down signals, which interact

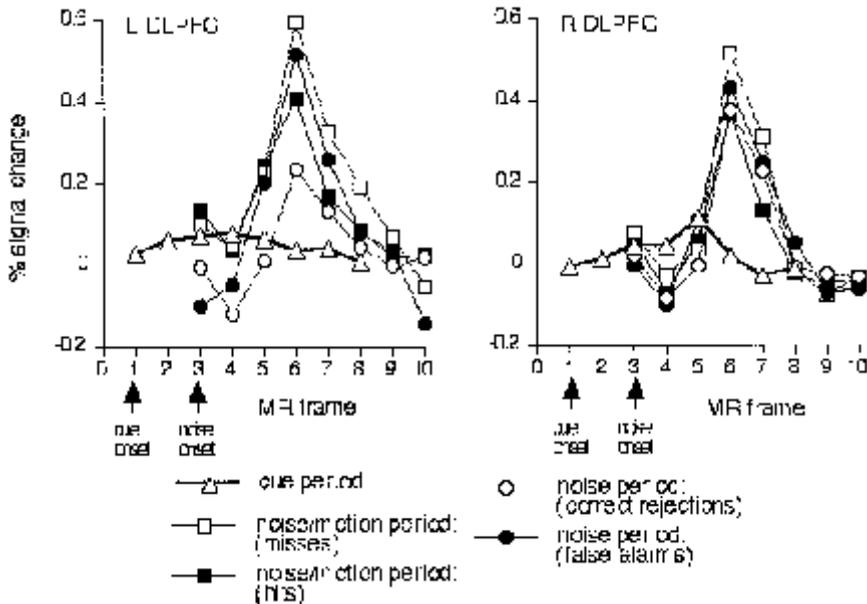


Figure 1.5

Group time courses in left and right dorsolateral prefrontal cortex (DLPFC) during the cue and test period. Test periods in which a motor response was executed (hits, false alarms) are indicated by filled symbols, and test periods in which a response was withheld (misses, correct rejections) are indicated by open symbols. No motor responses were made during cue periods. Time courses are based on a 3×3 voxel region of interest centered on the voxel yielding the maximum Z score in the region. BOLD responses are evident only during the test period, but are not contingent on whether a motion target was detected or a response was made.

in many motion-sensitive regions with incoming sensory stimulation, producing a stronger sensory response to incoming visual information. This (attentional) modulation leads, through unclear mechanisms, to a more accurate discrimination.

This experiment raises many new questions about visual expectations and top-down expectation signals. One would like to know more about the selectivity of such signals for direction of motion. In other experiments we found that the same regions were also active during the selection of target locations in the course of luminance detection tasks (Chelazzi et al., 2000; Corbetta et al, 2000). Other experiments have shown that attention to location and motion can have additive effects in modulating activity in the motion system (Beauchamp et al., 1997). Psychological results suggest that selection by a feature does not occur per se, but is mediated through location codes (Moore and Egeth, 1998; Shih and Sperling, 1996). Another important question is whether the network that maintains expectation signals for motion will change, depending on the format in which the cue is

presented. Is the posIPs activity reflecting some supramodal signal that will hold information independently of cue format, or will other regions come on-line by increasing the complexity of the transformation (e.g., DLPFC)? A final intriguing question is whether the level of top-down modulation will vary as a function of task demands. Is it possible to specifically preactivate different levels in the motion system, depending on whether the perceptual task relies on low- or high-level motion analysis (e.g., motion energy versus feature detection versus structure-from-motion; Sperling et al., chapter 10 in this volume)?

1.5 Relevance for Psychological and Neural Theories of Attention

The notion that top-down expectation signals preactivate sensory pathways is consistent with previously reviewed single-unit data during spatial attention (Luck et al., 1997) and match-to-sample tasks (Chelazzi et al., 1993). Tonic increases in the rate of baseline firing or delay activity related to the cue may represent putative neuronal correlates of the BOLD modulation described above. Pathway preactivation is also consistent with cognitive, computational, and neural models of higher vision. In Wolfe's guided search model (Wolfe et al., 1989), top-down expectation signals about the relevant feature bias activity in feature maps that represent the visual scene. In Tsotsos's model, a top-down biasing signal feeds back to feature-level maps, where it selectively activates winner-take-all competition circuits in order to emphasize relevant targets and deemphasize irrelevant distractors (Tsotsos et al., chapter 14 in this volume). In Desimone and Duncan's (1995) biased competition model, the bottom-up competition between objects in a visual scene is gradually resolved in favor of the object receiving a top-down biasing signal from frontal structures.

All these models implicitly assume that the top-down expectation signal has its primary effect on the relevant visual representation, that is, the one coding for the relevant object. The notion of pathway preactivation is also consistent with psychological theories previously discussed, although this experiment does not resolve between different accounts. The findings are certainly consistent with a signal enhancement mechanism, given the presence of preactivation in relevant pathways, and the fact that early levels in the motion system were modulated (MT+, lateral occipital). However, noise information was also coded in the same regions, so that we cannot rule out the possibility that expectation signals reflected noise suppression mechanisms. This issue could be solved by recording in the monkey from cortical columns that code for cued and uncued directions during the period of time following the symbolic cue. The cue's preactivation of irrelevant or relevant columns would be consistent, respectively, with a noise suppression or signal enhancement account.

Signal enhancement mechanisms are currently more strongly supported by the experimental evidence. Luck and colleagues (1997) observed in area V4 that the tonic increase

in baseline firing rate, a correlate of a top-down biasing signal for location, is maximal at the attended location and falls off sharply as one moves away from the focus of attention, even within the same receptive field. In other words, the top-down signal is stronger at (attended) locations processing signal than at (unattended) locations processing noise. As reviewed earlier, most imaging studies show attentional modulations in areas coding for the relevant attribute, with little or no modulation in areas coding for the irrelevant attribute. It will be important to show with event-related fMRI methods that expectation signals coding for one particular visual attribute (among many potentially relevant attributes) preactivate only relevant sensory pathways. In a recent imagery experiment, Kanwisher and O'Craven demonstrated the feature selectivity of some top-down signals (Kanwisher and O'Craven, 1998). They measured BOLD signal simultaneously in two extrastriate visual regions, specialized respectively for face and place/building processing, while subjects formed mental images of famous faces (e.g., Bill Clinton) or buildings (e.g., Eiffel Tower) after verbal expectation. The BOLD signal increased specifically in each area as a function of the expectation, and presumably the mental image, being formed. In other words, the BOLD signal increased in the *face* area when a famous person was named. The expectation signal triggered by the verbal expectation was therefore stimulus-specific and modulated the relevant neural representation.

1.6 Conclusions

Neuroimaging research on attention since 1990 has convincingly demonstrated the importance of visual expectations for visual processing. Extensive modulations throughout visual cortex have been demonstrated, using a variety of selection criteria and paradigms. Modulations occur predominantly in task-relevant pathways, and represent an interaction effect between top-down signals that encode/maintain visual expectations, and incoming sensory information. Recent advances in functional imaging methods have begun to provide a glimpse of the temporal dynamics of the processes involved in attentional control.

References

- Ball, K., and Sekuler, R. (1981). Cues reduce direction uncertainty and enhance motion detection. *Percept. Psychophys.* 30: 119–128.
- Beauchamp, M. S., Cox, R. W., and DeYoe, E. A. (1997). Graded effects of spatial and featural attention on human area MT and associated motion processing regions. *J. Neurophysiol.* 78: 516–520.
- Boynton, G. M., Engel, S. A., Glover, G. H., and Heeger, D. J. (1996). Linear systems analysis of fMRI in human V1. *J. Neurosci.* 16: 4207–4221.
- Brefczynski, J. A., and DeYoe, E. A. (1999). A physiological correlate of the “spotlight” of visual attention. *Nature Neurosci.* 2: 370–374.

- Britten, K. H., Shadlen, M. N., Newsome, W. T., and Movshon, J. A. (1993). The analysis of visual motion: A comparison of neuronal and psychophysical performance. *J. Neurosci.* 12: 4745–4765.
- Broadbent, D. (1982). Task combination and selective intake of information. *Acta Psychol.* 50: 253–290.
- Bruce, C. J., and Goldberg, M. E. (1985). Primate frontal eye fields. I. Single neurons discharging before saccades. *J. Neurophysiol.* 53: 603–635.
- Buchel, C., Josephs, O., Rees, G., Turner, R., Frith, C., and Friston, K. (1998). The functional anatomy of attention to visual motion. *Brain* 121: 1281–1294.
- Buckner, R. L., Bandettini, P. A., O’Craven, K. M., Savoy, R. L., Petersen, S. E., Raichle, M. E., and Rosen, B. R. (1997). Detection of transient and distributed cortical activation during single trials of a cognitive task using fMRI. *Proc. Nat. Acad. Sci. USA* 93: 14878–14883.
- Chelazzi, L., Miller, E. K., Duncan, J., and Desimone, R. (1993). A neural basis for visual search in inferior temporal cortex. *Nature* 363: 345–347.
- Colby, C. L., Duhamel, J. R., and Goldberg, M. E. (1996). Visual, presaccadic, and cognitive activation of single neurons in monkey lateral intraparietal area. *J. Neurophysiol.* 76: 2841–2852.
- Corbetta, M., Miezin, F. M., Dobmeyer, S., Shulman, G. L., and Petersen, S. E. (1991). Selective and divided attention during visual discriminations of shape, color, and speed: Functional anatomy by positron emission tomography. *J. Neurosci.* 11: 2383–2402.
- Dale, A. M., and Buckner, R. L. (1997). Selective averaging of rapidly presented individual trials using fMRI. *Hum. Brain Mapping* 5: 329–340.
- Desimone, R., and Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annu. Rev. Neurosci.* 18: 193–222.
- Deutsch, J., and Deutsch, D. (1963). Attention: Some theoretical considerations. *Psychol. Rev.* 70: 80–90.
- Duncan, J. (1980). The locus of interference in the perception of simultaneous stimuli. *Psychol. Rev.* 87: 272–300.
- Dupont, P., Orban, G. A., Vogels, R., Bormans, G., Nuyts, J., Schiepers, C., De Roo, M., and Mortelmans, L. (1993). Different perceptual tasks performed with the same visual stimulus attribute activate different regions of the human brain: A positron emission tomography study. *Proc. Nat. Acad. Sci. USA* 90: 10927–10931.
- Engel, A. K., König, P., and Singer, W. (1991). Direct physiological evidence for scene segmentation by temporal coding. *Proc. Nat. Acad. Sci. USA* 88: 9136–9140.
- Eriksen, C. W., and Hoffman, J. E. (1972). Temporal and spatial characteristics of selective encoding from visual displays. *Percept. Psychophys.* 12: 201–204.
- Friston, K. J., Josephs, O., Rees, G., and Turner, R. (1998). Nonlinear event-related responses in fMRI. *Magn. Reson. Med.* 39: 41–52.
- Funahashi, S., Chafee, M. V., and Goldman-Rakic, P. S. (1993). Prefrontal neuronal activity in rhesus monkeys performing a delayed anti-saccade task. *Nature* 365: 753–756.
- Fuster, J. M. (1973). Unit activity in prefrontal cortex during delayed-response performance: Neuronal correlates of transient memory. *J. Neurophysiol.* 36: 61–78.
- Gandhi, S. P., Heeger, D. J., and Boynton, G. M. (1999). Spatial attention affects brain activity in human primary visual cortex. *Proc. Nat. Acad. Sci. USA* 96: 3314–3319.
- Gilbert, C. D. (1996). Plasticity in visual perception and physiology. *Curr. Opin. Neurobiol.* 6: 269–274.
- Haenny, P. E., Maunsell, J. H. R., and Schiller, P. H. (1988). State dependent activity in monkey visual cortex. II. Retinal and extraretinal factors in V4. *Exp. Brain Res.* 69: 245–259.
- Hawkins, H. L., Hillyard, S. A., Luck, S. J., Mouloua, M., Downing, C. J., and Woodward, D. P. (1990). Visual attention modulates signal detectability. *J. Exp. Psychol. Hum. Percept. Perf.* 16: 802–811.
- Haxby, J. V., Horowitz, B., Ungerleider, L. G., Maisog, J. M., Pietrini, P., and Grady, C. L. (1994). The functional organization of human extrastriate cortex: A PET–rCBF study of selective attention to faces and locations. *J. Neurosci.* 14: 6336–6353.

- Hillyard, S. A., and Picton, T. W. (1987). Electrophysiology of cognition. In F. Plum, V. B. Mountcastle, and S. T. Geiger (eds.), *The handbook of physiology*. Sect. 1, *The nervous system*, vol. 5, *Higher functions of the brain*, pt. 2 (pp. 519–584). Bethesda, MD: American Physiological Society.
- Hillyard, S. A., Vogel, E. K., and Luck, S. J. (1998). Sensory gain control (amplification) as a mechanism of selective attention: Electrophysiological and neuroimaging evidence. *Philos. Trans. R. Soc. London B* 353: 1257–1270.
- Ito, M., Westheimer, G., and Gilbert, C. D. (1998). Attention and perceptual learning modulate contextual influences on visual perception. *Neuron* 20: 1191–1197.
- Kanwisher, N., and O’Craven, K. M. (1998). Extrastriate activity during visual imagery is stimulus specific. *Soc. Neurosci. Abstr.* 24: 530.
- Karni, A., Meyer, G., Jezard, P., Adams, M. M., Turner, R., and Ungerleider, L. G. (1995). Functional MRI evidence for adult motor cortex plasticity during motor skill learning. *Nature* 377: 155–158.
- Kastner, S., De Weerd, P., Desimone, R., and Ungerleider, L. G. (1998). Mechanisms of directed attention in the human extrastriate cortex as revealed by functional MRI. *Science* 282: 108–111.
- Knierim, J. J., and Van Essen, D. C. (1992). Neuronal responses to static texture patterns in area V1 of the alert macaque monkey. *J. Neurophysiol.* 67: 961–980.
- Logothetis, N. K., Guggenberger, H., Peled, S., and Pauls, J. (1999). Functional imaging of the monkey brain. *Nature Neurosci.* 2: 555–562.
- Luck, S. J., Chelazzi, L., Hillyard, S. A., and Desimone, R. (1997). Neuronal mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *J. Neurophysiol.* 77: 24–42.
- Mangun, G. R., Hillyard, S. A., and Luck, S. J. (1993). Electrocortical substrates of visual selective attention. In D. Meyer and S. Kornblum (eds.), *Attention and performance XIV* (pp. 219–243). Cambridge, MA: MIT Press.
- Martinez, A., Anllo-Vento, L., Sereno, M. I., Frank, L. R., Buxton, R. B., Dubowitz, D. J., Wong, E. C., Hinrichs, H., Heinze, H. J., and Hillyard, S. A. (1999). Involvement of striate and extrastriate visual cortical areas in spatial attention. *Nature Neurosci.* 2: 364–369.
- Maunsell, J. H. R. (1995). The brain’s visual world: Representation of visual targets in cerebral cortex. *Science* 270: 764–769.
- Miller, E. K., Erickson, C. A., and Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *J. Neurosci.* 16: 5154–5167.
- Moore, C. M., and Egeth, H. (1998). How does feature-based attention affect visual processing? *J. Exp. Psychol. Hum. Percep. Perf.* 24: 1296–1310.
- Moran, A. J., and Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science* 229: 782–784.
- Motter, B. C. (1993). Focal attention produces spatially selective processing in visual cortical areas V1, V2, and V4 in the presence of competing stimuli. *J. Neurophysiol.* 70: 909–919.
- Motter, B. C. (1994a). Neural correlates of attentive selection for color or luminance in extrastriate area V4. *J. Neurosci.* 14: 2178–2189.
- Motter, B. C. (1994b). Neural correlates of feature selective memory and pop-out in extrastriate area V4. *J. Neurosci.* 14: 2190–2199.
- Mountcastle, V. B., Motter, B. C., Steinmetz, M. A., and Sestokas, A. K. (1987). Common and differential effects of attentive fixation on the excitability of parietal and prestriate (V4) cortical visual neurons in the macaque monkey. *J. Neurosci.* 7: 2239–2255.
- Nystrom, L. E., Delgado, M. R., Sabb, F. W., Noll, D. C., and Cohen, J. (1998). Dynamics of fMRI: Broca’s area activation reflects independent effects of duration and intensity of working memory processes. *NeuroImage* 7: 7.
- O’Craven, K. M., Savoy, R. L., and Rosen, B. R. (1995). Attention modulates fMRI activation in human MT/MST. *Soc. Neurosci. Abstr.* 21: 1988.
- Ollinger, J. M., Corbetta, M., Petersen, S. E., and Shulman, G. L. (1998). Separating events within a trial with event-related fMRI. *Soc. Neurosci. Abstr.* 24: 1178.

- Olshausen, B. A., Anderson, C. H., and Van Essen, D. C. (1993). A neurobiological model of visual attention and invariant pattern recognition based on dynamic routing of information. *J. Neurosci.* 13: 4700–4719.
- Pashler, H. E. (1998). *The psychology of attention*. Cambridge, MA: MIT Press.
- Posner, M. I. (1980). Orienting of attention. *Q. J. Exp. Psychol.* 32: 3–25.
- Raichle, M. E. (1994). Visualizing the mind. *Sci. Am.* 270: 58–64.
- Rosen, B. R., Buckner, R. L., and Dale, A. M. (1998). Event related fMRI: Past, present, and future. *Proc. Nat. Acad. Sci. USA* 95: 773–780.
- Shih, S., and Sperling, G. (1996). Is there feature-based attentional selection in visual search? *J. Exp. Psychol. Hum. Percep. Perf.* 22: 758–779.
- Shulman, G. L., Ollinger, J. M., Akbudak, E., Conturo, T. E., Snyder, A. Z., Petersen, S. E., and Corbetta, M. (1999). Areas involved in encoding and applying directional expectations to moving objects. *J. Neurosci.* 19: 9480–9496.
- Somers, D. C., Dale, A. M., Seiffert, A. E., and Tootell, R. B. H. (1999). Functional MRI reveals spatially specific attentional modulation in human primary visual cortex. *Proc. Nat. Acad. Sci. USA* 16: 1663–1668.
- Stefanacci, L., Reber, P., Costanza, J., Way, E., Buxton, R., Zola, S., Squire, L., and Albright, T. (1998). fMRI of monkey visual cortex. *Neuron* 20: 1051–1057.
- Tootell, R. B. H., Dale, A. M., Sereno, M. I., and Malach, R. (1996). New images from human visual cortex. *Trends Neurosci.* 19: 481–489.
- Tootell, R. B. H., Hadjikhani, N., Hall, E. K., Marrett, S., Vanduffel, W., Vaughan, J. T., and Dale, A. M. (1998). The retinotopy of visual spatial attention. *Neuron* 21: 1409–1422.
- Treisman, A. M. (1969). Strategies and models of selective attention. *Psychol. Rev.* 76: 282–299.
- Treue, S., and Maunsell, J. H. R. (1996). Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature* 382: 539–541.
- Ts'o, D. Y., Frostig, R. D., Lieke, E. E., and Grinvald, A. (1990). Functional organization of primate visual cortex revealed by high resolution optical imaging. *Science* 27: 417–420.
- Van Essen, D. C., and DeYoe, E. A. (1995). Concurrent processing in the primate visual system. In M. S. Gazzaniga (ed.), *The cognitive neurosciences* (pp. 383–400). Cambridge, MA: MIT Press.
- Vandenberghe, R., Duncan, J., Dupont, P., Ward, R., Poline, J. B., Bormans, G., Michiels, J., Mortelmans, L., and Orban, G. A. (1997). Attention to one or two features in left and right visual field: A positron emission tomography study. *J. Neurosci.* 17: 3739–3750.
- Vandenberghe, R., Dupont, P., Debruyn, B., Bormans, G., Michiels, J., Mortelmans, L., and Orban, G. A. (1996). The influence of stimulus location on the brain activation pattern in detection and orientation discrimination: A PET study of visual attention. *Brain* 119: 1263–1276.
- Wojciulik, E., Kanwisher, N., and Driver, J. (1998). Covert visual attention modulates face-specific activity in the human fusiform gyrus: fMRI study. *J. Neurophysiol.* 79: 1574–1578.
- Woldorff, M. G., Fox, P. T., Matzke, M., Lancaster, J. L., Veeraswamy, S., Zamarripa, F., Seabolt, M., Glass, T., Gao, J. H., Martin, C. C., and Jerabek, P. (1997). Retinotopic organization of early visual spatial attention effects as revealed by PET and ERPs. *Hum. Brain Mapping* 5: 280–286.
- Wolfe, J. M., Cave, K. R., and Franzel, S. L. (1989). Guided search: An alternative to the feature integration model for visual search. *J. Exp. Psychol. Hum. Percep. Perf.* 15: 419–433.
- Wurtz, R. H., Goldberg, M. E., and Robinson, D. L. (1980). Behavioral modulation of visual responses in the monkey: Stimulus selection for attention and movement. *Prog. Psychobiol. Physiol. Psychol.* 9: 43–83.
- Wurtz, R. H., Richmond, B. J., and Newsome, W. T. (1984). Modulation of cortical visual processing by attention, perception, and movement. In G. M. Edelman, W. E. Gall, and W. M. Cowan (eds.), *Dynamic aspects of neocortical functions* (pp. 195–217). New York: Wiley.
- Zarahn, E., Aguirre, G. K., and D'Esposito, M. (1997). Empirical analyses of BOLD fMRI statistics. I. Spatially unsmoothed data collected under null-hypothesis conditions. *NeuroImage* 5: 179–197.