CHAPTER 4

A MULTI-LEVEL ACCOUNT OF SELECTIVE ATTENTION

JOHN T. SERENCES AND SABINE KASTNER

INTRODUCTION

NEURAL information processing systems must overcome a series of bottlenecks that interrupt the sequence of events between sensory input and motor output. First, the sensory neurons that encode external stimuli are noisy. As a result, an identical stimulus will evoke a slightly different response pattern each time it is presented, and this instability can place limits on the amount of information that neurons can relay about stimulus features (Pouget, Dayan, and Zemel 2003; Seung and Sompolinsky 1993). Second, multiple items in the visual field compete for representation, and this competition must be resolved so that the most behaviourally relevant sensory stimuli are represented and allowed to guide goal-directed behaviour. Third, there are limits on the number of responses that can be selected simultaneously (Pashler 1994), and obvious limits on the number of simultaneous motor plans that can be executed. All of these factors restrict the speed and accuracy of sensory–response mappings, and at each stage, relevant stimuli must be prioritized over irrelevant distractors to ensure the successful completion of behavioural goals. The ability to prioritize relevant stimuli is generally referred to as *selective attention*, where the prefix *selective* is intentionally used to distinguish the term from changes in general arousal or states of consciousness.

There are three general mechanisms that support selection attention: signal enhancement, external and internal noise suppression, and selective read-out. The first two mechanisms influence selection by directly mediating the information content of neural representations as sensory signals are relayed and transformed across the processing hierarchy. Selective read-out is conceptually different, as it does not directly modify the quality of sensory representations. Instead, selective read-out influences the efficiency with which sensory information is used to inform decisions about the contents of the sensory array, either in the absence of or in combination with signal enhancement and noise suppression.

In this chapter, we focus on two general questions that have emerged over the last 50 years: *where* and *when* do these mechanisms of selective attention operate to enable efficient information processing? Such questions about the locus of selection are rooted in the classic debate between theories of early and late selection that have been actively pursued in the fields of cognitive psychology and cognitive neuroscience for decades. Early selection theories hold that attention filters sensory inputs based on simple low-level features (e.g. pitch of a sound, orientation of a line), whereas late selection theories hold that filtering occurs only after all basic features have been analysed and stimuli are semantically labelled. More recently, this early/late dichotomy has been subjected to increasingly sophisticated neuroscientific techniques that reveal signatures of selection in large-scale neural networks including the thalamus, sensory, and higher-order cortex. At different points in time and at different levels of the cortical hierarchy, selection can be implemented via signal enhancement, internal and external noise suppression, and the selective read-out of sensory signals. Moreover, these mechanisms operate in a complementary manner to facilitate information processing, with the relative contribution of each determined by behavioural demands and by the nature of the stimulus array.

We begin with a brief history of the early versus late selection debate, and then review current knowledge about *where* and *when* signal enhancement, noise suppression, and selective read-out exert their influence on information processing.

A Historical Perspective on Early Versus Late Selection Theories

One of the first clearly articulated theories concerning the locus of attentional selection was the filter model of Broadbent (Broadbent 1958). He posited that incoming stimuli were analysed to the level of basic features (e.g. pitch, location). Based on this information, one item at a time was then selected and brought into awareness. Critically, given the breakdown of information into basic components assumed in this model, selection must happen very early in the processing stream, before the level of semantic analysis. Thus, models of this type are termed *early selection* theories. Broadbent was primarily influenced by data from dichotic listening tasks, where subjects were presented with separate auditory input in each ear. In these classical studies, subjects were typically asked to shadow the input from one ear by repeating out loud. While this task is quite attentionally engaging, subjects have little problem selecting one of the two auditory streams to shadow based on low-level features such as the pitch of the speaker's voice or the ear of origin (location). Moreover, subjects typically report having little subjective awareness of the information presented in the unattended auditory channel, as reflected quantitatively in poor recognition memory performance (Cherry 1953; Moray 1959; Neisser and Becklen 1975). Nevertheless, subjects were still able to discriminate low-level stimulus attributes such as pitch in the unattended channel, supporting

Broadbent's notion that selection must be early, as these basic attributes appeared to be all that was retained from the unattended auditory stream.

Not long after Broadbent's seminal book, Moray (1959) demonstrated that selection was not always implemented by an early filtering mechanism, as he noted that about one-third of subjects detected their own name when it was inserted in the unattended stream, despite a general lack of awareness of the overall content of the message. In a related finding, Treisman (Treisman 1960) presented subjects with two coherent sentences, one in each ear, and subjects were instructed to shadow the input from one of the channels. When the flow of the prose from the shadowed channel changed to the other ear, subjects would often continue to shadow the same sentence, even though the words were now being presented in the to-be-ignored channel. This tendency to track a sentence from the attended to the unattended channel based on syntactic structure and semantic content suggests that the message in the unattended ear was processed to a far more advanced level of analysis than predicted by the original inception of Broadbent's early selection theory.

To account for the semantic processing of the unattended channel, Anthony and Diane Deutsch (Deutsch and Deutsch 1963) formalized a *late selection* theory of attentional selection. In stark contrast to Broadbent's early filter theory, the late selection account holds that all stimuli are analysed to the level of semantic description. Moreover, the mechanism that supports semantic analysis is not capacity limited and can thus process all perceptual inputs in parallel. After this analysis is complete, attention operates by selecting a subset of the items for further processing, such as encoding into working memory (Duncan 1980). Importantly, Deutsch and Deutsch did not imply that subjects necessarily had conscious access to these early semantic descriptions. Rather, awareness only occurred after the capacity-limited process of attentional selection, accounting for the commonly held notion that observers are only aware of a small subset of potential stimuli at any given moment in time.

After these two opposite views on the locus of selection had been established, a good deal of research over the next 30 years focused on attempting to critically discriminate between the competing models. While studies on dichotic listening remained prominent (e.g. Corteen and Wook 1972), many researchers turned to investigating early and late selection in the visual domain. The prototypical approach was to infer the locus of selection based on behavioural measures such as priming that reflect the extent to which unattended stimuli impacted performance. Instances in which unattended stimuli did not influence performance were taken as support for early selection. Conversely, the influence of unattended stimuli on performance was interpreted as evidence in favour of late selection. For instance, Eriksen and Eriksen (Eriksen and Eriksen 1974) had subjects report the identity of a target letter briefly presented in an attended location indicated by a pre-cue presented at the start of each trial. One of the two possible target letters was mapped to a button press with the index finger, and the other target letter was mapped to a button press with the middle finger. The target letter was flanked on either side by additional letters that were either mapped to the same response required by the target (termed *compatible flankers*) or to the opposite response (termed *incompatible flankers*). The main observation was that response times (RT) to the targets were slower when the target was accompanied by incompatible flankers as opposed to compatible flankers. The effect of flanker identity on RT supports the late selection tenet that items outside the locus of attention are processed at least to the level of semantic description despite the fact that subjects had perfect foreknowledge of the target's location and could presumably focus attention well in advance.

However, evidence of this sort is not undisputed. First, proponents of early selection can always assert that the behavioural tasks did not adequately engage the subject's attention, thus 'leakage' through an early perceptual filter might occur. Indeed, if attentional cues are made salient enough, then irrelevant flankers do not have a measurable effect on performance, as shown by Egeth and colleagues (Francolini and Egeth 1980; Yantis and Johnston 1990). Such observations undermine the strong form of the late selection account because late selection models clearly predict an effect of incompatible distractors, regardless of the attentional focus. However, it appears that these claims against late selection rest on the assumption that the dependent measure is sensitive to a subtle influence of distractors. Indeed, using a different behavioural measure, Tipper and Driver (Tipper and Driver 1988) found evidence that the distractors used by Egeth and colleagues were in fact processed to the level of semantic content, supporting a late selection account. Moreover, the view that a semantic description was formed and then suppressed prior to the response stage can always be asserted by supporters of late selection to countermand data consistent with early selection. The main conclusion of these behavioural studies, and many others like them, is that behavioural evidence in support of either early or late selection can be obtained using very similar experimental paradigms. Therefore, in this particular area of cognitive psychology, physiological metrics have had a great deal of influence, as they can provide more direct insight into the temporal dynamics of signal enhancement, noise suppression, and the level to which sensory information gets processed.

As with the behavioural studies reviewed above, the majority of neuroimaging and electrophysiology studies investigating the neural correlates underlying attentional selection have focused on the visual rather than the auditory domain. Thus, we will limit our discussion to studies on the visual system, first discussing issues pertaining to early and late selection in terms of 'where' within the visual processing hierarchy selection plays out and then asking 'when' these modulations occur.

THE 'WHERE' OF SELECTION VIA SIGNAL ENHANCEMENT: NEURAL EVIDENCE FROM MONKEY PHYSIOLOGY AND HUMAN NEUROIMAGING

The first important question that emerges from the 'early' vs. 'late' selection debate that we will consider at the neural level is: where in the visual pathway does selective attention *first* affect neural processing? Early and late selection accounts appear to make straightforward predictions with respect to this question. An early selection theory would posit that selective attention should influence early stages of processing including primary visual cortex (striate cortex, or V1), where visual information is filtered based on basic features such as colour, motion, and orientation. In contrast, a late selection theory would predict no modulation at early processing stages, as information should be faithfully represented regardless of task demands (at least) up to stages where the visual information interfaces with other cognitive domains such as memory or action context.

The modulation of neural responses by spatially selective attention was originally demonstrated in the extrastriate cortex. In a seminal study, Moran and Desimone (Moran and Desimone 1985) recorded the activity of neurons in area V4 from monkeys trained to direct attention to a spatial location within a neuron's receptive field (RF). The target stimulus was either shown alone in the neuron's RF or in the presence of a nearby distractor. Attending to the target enhanced neural responses, but only when a competing stimulus was presented in the same RF as the target. These findings demonstrate that selective attention can gate the processing of behaviourally relevant information by effectively constricting the RF around the selected stimulus such that the distractor has little or no impact on the firing rate of the cell. Similar effects of attentional enhancement in area V4 have been demonstrated by several other groups (e.g. Connor, Preddie, Gallant, and Van Essen 1997; Haenny and Schiller 1988; McAdams and Maunsell 1999; Motter 1993), and have been found in several other extrastriate and parietal areas including V2 (Luck, Chelazzi, Hillyard, and Desimone 1997; Motter 1993), MT (e.g. Treue and Maunsell 1996) and LIP (e.g. Gottlieb, Kusunoki, and Goldberg 1998).

Despite the widespread and robust observation of attentional modulation in many areas of visual cortex, it was initially unclear if attention modulated activity in primary visual cortex. Several studies were unable to demonstrate such modulation (e.g. Luck et al. 1997), while others found relatively weak, but reliable effects (e.g. Motter 1993). While it was not clear what factors best accounted for these discrepant findings, there was an overall impression that attentional modulation at the earliest cortical processing stages may be highly dependent on task-related factors or the need to integrate information from beyond the classic RF (e.g. Ito and Gilbert 1999; Roelfsema, Lamme, and Spekreijse 1998). In either case, the few studies that recorded attention effects from different areas in the same animals (Cook and Maunsell 2002a, 2002b; Luck et al. 1997) showed that the magnitude of attentional modulation was stronger in more anterior extrastriate areas (e.g. V4) compared to more posterior areas such as V2. This graded effect was taken to suggest that attentional effects at earlier stages were caused by reactivation from higher-order extrastriate areas through cortico–cortical feedback connections, although a direct demonstration of this link has not been reported.

Thus, evidence from these initial physiology studies does not unequivocally support either early or late selection models. The finding of attentional modulation in relatively early cortical areas such as V2 and possibly even V1 is consistent with early selection accounts. However, the inconsistent observation of attention effects in V1, coupled with the possibility that these effects primarily reflect reactivation from higher-order areas, renders the evidence inconclusive. Thus, single unit recording data that emerged by the late 1990s did not clearly support either alternative, leaving the question far from being settled.

In the mid and late 1990s, the advent of functional magnetic resonance imaging (fMRI) enabled detailed studies of the human visual system and offered several advantages over previous approaches. In particular, human subjects can perform a wider range of tasks compared to non-human primates and activation changes tied to task manipulations can be examined across the entire brain (as opposed to monitoring single-unit activity in a single area). Soon, investigators had settled lingering questions about whether attention could modulate responses at the earliest anatomical levels of processing by showing response modulation in V1 with moving (Gandhi, Heeger, and Boynton 1999; Somers, Dale, Seiffert, and Tootell 1999) and stationary stimuli (Martinez et al. 1999). Even stronger evidence in support of early selection came from the subsequent demonstration of attention effects in the thalamus, i.e. the lateral geniculate nucleus (LGN) (O'Connor, Fukui, Pinsk, and Kastner 2002). The LGN is the thalamic component in the retinocortical projection and it is the first neural structure within the visual processing hierarchy that can be modulated by feedback input (via afferent input from the thalamic reticular nucleus, striate cortex, and the brainstem). Interestingly, attention effects in the LGN were found to be stronger than in area V1, more on the order of extrastriate areas such as V4 (O'Connor et al. 2002). Thus, these findings not only challenge the notion that cognitive processing primarily involves cortical networks, but also the notion that attentional modulation in early sensory areas can be explained solely by cortico–cortical re-entrant signals from later stages of the hierarchy, as such an account would have predicted smaller effects of attention in the LGN than in V1. Instead, fMRI signals may reflect the summed modulatory feedback that a given area receives from cortical, thalamic, and brainstem sources, as opposed to just local spiking activity. If this is indeed the case, then larger attentional modulation in LGN might be related to the larger number of afferent inputs that this region receives compared to primary visual cortex. Thus, fMRI evidence gathered over the last 15 years provides compelling support for anatomically early attention effects by demonstrating modulation in V1 and the LGN, the earliest stages of visual processing that receive feedback influences from other sources (see also chapters by Beck and Kastner, and Saalmann and Kastner).

The ability of fMRI measurements to index attentional modulation across the entire brain also provides key insights into how each area is driven by 'sensory-driven' or 'bottom-up' versus 'cognitively driven' or 'top-down' influences. For instance, sensory-driven population responses during passive viewing follow a characteristic pattern when measured with fMRI: the strongest responses are typically observed in early visual areas such as V1 and progressively weaker responses are observed at successively later stages of the hierarchy (see Fig. 4.1a). For example, 90% of the maximum response can be evoked in V1 by a simple visual stimulus, whereas the same input might evoke a modulation of only 10% in higher-order areas of parietal and frontal cortex (e.g. Treue 2003). On the other hand, the response modulation associated with deploying

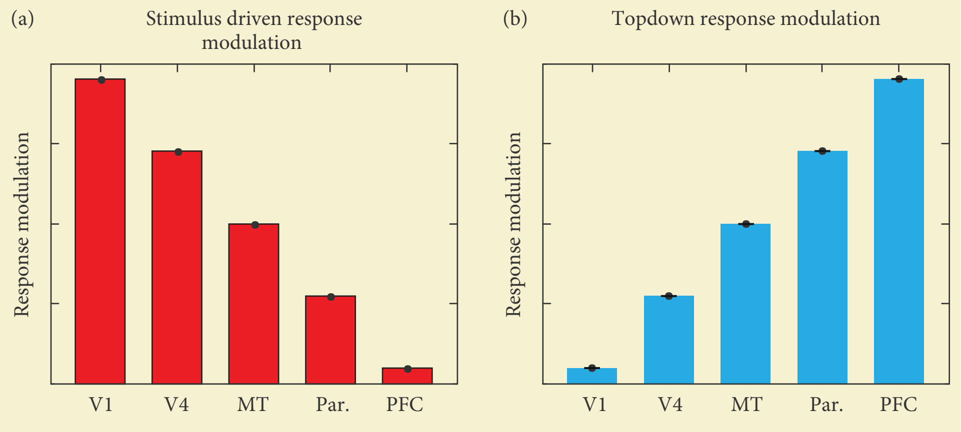

FIGURE 4.1 Schematic showing the relative magnitude of stimulus-driven and top-down modulation across the cortical hierarchy from occipital to parietal (Par.) to pre-frontal cortex (PFC). (a) The magnitude of sensory modulation gradually decreases. (b) In contrast, the magnitude of top-down modulations is largest in PFC and gradually decreases in earlier visual areas.

top-down attention to a fixed sensory stimulus exhibits the opposite pattern: a 90% increase in activation might be expected in areas of parietal and frontal cortex, whereas a modulation of 10% is expected in early visual cortex (see Fig. 4.1b).

Based on the neural evidence of attentional modulation at early processing stages, the preponderance of evidence seems to favour early selection as opposed to late selection accounts. However, several caveats should be noted. First, none of the studies reviewed thus far addressed the question of *what* was selected or what type of information was modulated at each processing stage. This is a critical issue, as the observation of an attentional modulation at a specific stage of the visual hierarchy does not speak to the issue of what kind of information is targeted by attention. For example, an increase in activation levels within a neural population that is thought to encode a certain feature—such as motion in area MT—does not necessarily mean that selective attention operates on the feature 'motion'. This is particularly true when neural effects are inferred based on fMRI measurements, as the blood oxygen level dependent (BOLD) signal reflects the aggregate activity across all neurons in a region and is generally insensitive to feature selective changes in population response profiles (see Serences and Saproo 2012). Thus, any attentional modulation that is measured within a given region may be targeted on any type of information processed in the area, or may simply reflect a general increase in the activity of all neurons within a region in a non-selective manner (i.e. analogous to a change in general arousal). Such effects may thus translate into a scaling of responses in a region without actually influencing the selectivity of the population response or the amount of stimulus-specific information that is encoded about relevant sensory stimuli. In addition, even though neuroimaging methods excel at pinpointing the anatomical locus of selection, fMRI does not have the temporal resolution to conclusively distinguish between modulations of afferent signals and modulations related to later re-entrant feedback. In contrast, electrophysiology studies that examine the temporal dynamics of the selection process have been able to shed more light on these issues by exploiting temporally precise markers that are thought to index the depth of information processing.

The 'When' of Selection via Signal Enhancement: the Temporal Dynamics of Selective Attention

As noted in the last section, the evidence is now compelling that attention can modulate neural activity very early in the visual processing stream (e.g. in LGN), consistent with early selection accounts. However, this is a necessary but not sufficient condition to support the early filtering of sensory information, as early selection also requires that the modulations occur early in *time*, before extensive processing has taken place at later stages. While still a matter of debate, studies that use human subjects and event-related potentials (ERPs) provide some of the most diagnostic information regarding the timing of attention effects.

The ERPs elicited by visual stimuli are generally decomposed into several components, the C1 (starting about 50–90 ms post-stimulus), the P1 (80–130 ms), and the N1 (140–200 ms). In a typical experimental paradigm, subjects are told to attend to a specific location in the visual field and a stimulus is either flashed at the attended location or at another, unattended, location. The classic finding is that the P1 and N1 components evoked by an attended stimulus are larger than the same components evoked by an unattended stimulus (see Mangun, Hillyard, and Luck 1993). Moreover, recent studies suggest that the earliest ERP component—the C1—can also be modulated by attentional factors, despite its onset as early as 50 ms after the presentation of a stimulus. Similarly early modulations have been shown in studies examining spatial attention (Kelly, Gomez-Ramirez, and Foxe 2008; Slotnick, Hopfinger, Klein, and Sutter 2002; Zhang, Zhaoping, Zhou, and Fang 2012—although see Martinez et al. 1999; Noesselt et al. 2002), auditory–visual coupling (Van der Burg, Talsma, Olivers, Hickey, and Theeuwes 2011), and in studies that examined the exogenous (or stimulus-driven) capture of attention to one of two superimposed visual surfaces (Khoe, Mitchell, Reynolds, and Hillyard 2005). Moreover, the C1 can be reasonably localized to V1 because the polarity of the response reverses when stimuli are presented in the upper and lower visual fields, consistent with the known anatomical layout and retinotopy of V1. Thus, the early onset of attention-related increases in response amplitude has led many researchers to conclude that attention modulates the afferent volley of sensory responses as they ascend the cortical hierarchy (Hillyard and Anllo-Vento 1998; Hillyard, Vogel, and Luck 1998). This evidence provides solid support for early selection, because strong forms of late selection cannot accommodate the differential processing of attended stimulus features during the initial feedforward sweep of sensory activity.

Taken together, the observation of enhanced early ERP responses suggests that attention mimics the effect of a physical increase in the salience of the sensory input (Hillyard and Anllo-Vento 1998; Hillyard et al. 1998). Importantly, early gain signals have also been associated with attention to basic features such as colour (Zhang and Luck 2009), which is consistent with Broadbent's original notion that early selection operates on basic stimulus features, not just on spatial locations. Importantly, these early amplitude modulations of ERP waveforms clearly violate assumptions made by pure late selection theories of attention that posit equivalent sensory analysis of all items present in the visual (or auditory) scene up to a relatively advanced levels of analysis. In addition, these early modulations are consistent with a recent line of work by Carrasco and co-workers, who have shown that attention actually increases the perceived salience of basic feature properties so that contrast becomes brighter, motion more salient, and so on (Carrasco 2009; Carrasco, Ling, and Read 2004; Carrasco and McElree 2001; Fuller, Park, and Carrasco 2009; Liu, Abrams, and Carrasco 2009; Liu, Fuller, and Carrasco 2006; see chapter by Carrasco; see also: Anton-Erxleben, Abrams, and Carrasco 2011; Carrasco, Fuller, and Ling 2008; Prinzmetal, Long, and Leonhardt 2008; Schneider 2006, 2011; Schneider and Komlos 2008). Collectively, these results support the intuitive notion that modulating the gain of responses in early visual areas should have a corresponding impact on the perceptual experience of attended features, contributing not only to increases in behavioural performance, but also to qualitative shifts in phenomenal experience.

Even though these results provide strong support for the feasibility of early selection accounts, there are several caveats. First, in order to directly compare responses evoked by attended and unattended stimuli, one must assume that the stimuli are processed in an identical manner except for the presence/absence of attentional gain. Presumably, if fundamentally different neural mechanisms were responsible for the observed differences in the response to attended and unattended stimuli, then the voltage distribution across the scalp would differ across these conditions. However, the voltage distribution across the scalp is nearly identical for attended and unattended stimuli (e.g. Mangun et al. 1993), suggesting that attention-related differences in ERP amplitude are indeed driven by changes in sensory gain and not by changes in the nature of the processing that is being carried out. Second, early sensory gain models predict that any amplitude and/or phase shifts due to attention should mirror those produced by actual changes in the physical attributes of the stimulus. In line with the predictions of early selection models, increasing stimulus brightness produces an increase in P1 and N1 amplitude that mirrors increases due to attention (particularly in the case of the P1; Johannes, Munte, Heinze, and Mangun 1995; Wijers, Lange, Mulder, and Mulder 1997).

Another key piece of evidence on the timing of attentional modulation comes from monkey single cell recording studies, which have demonstrated two patterns of neural activity that are consistent with early selection: an increase in spontaneous activity due to allocation of spatial attention and before the onset of a stimulus, and the rapid onset of a modulatory attention effect on visually evoked activity, presumably before recurrent feedback can influence the magnitude of spiking. Luck et al. (1997) were able to show both of these effects in single neurons recorded from the same animals (see also Kastner, Pinsk, De Weerd, Desimone, and Ungerleider 1999). In their paradigm, a monkey was trained to covertly attend to one of two spatial locations and to detect a pre-specified target at that location. In one condition, a single stimulus was presented within the RF of a V4 neuron and another stimulus was presented in the mirror symmetrical position across the vertical meridian (outside of the neuron's RF). When attention was directed inside the neuron's RF before the stimulus display was presented, there was an increase in the spiking rate of the neuron relative to when attention was directed to the other location outside of the RF. This divergence in activity, or *baseline shift*, was found during the 500 ms epochs before the onset of the stimulus display (see also Kastner et al. 1999; Ress, Backus, and Heeger 2000; Serences, Yantis, Culberson, and Awh 2004). Therefore, the observed modulation in spiking rate was not due to exogenous sensory stimulation. A second condition in the paradigm of Luck et al. (1997) evaluated the effects of attention on the post-stimulus activity of neurons in V4. To achieve maximum post-stimulus attention effects, both the target and the distractor were presented within the RF of a single V4 neuron (as in Moran and Desimone 1985). The monkey attended to one of the two possible stimulus locations, and spiking activity to the onset of a stimulus was recorded as a function of the locus of attention. Post-stimulus histograms revealed a clear separation of activity approximately 60 ms after stimulus onset such that attended stimuli evoked a larger response compared to unattended stimuli (see also Buffalo, Fries, Landman, Liang, and Desimone 2010). Similar results were also observed for single neurons in area V2, and are consistent with the similarly early modulation of the C1 ERP component measured from human subjects. Interestingly, however, no baseline shifts were observed in V1, perhaps because the RFs were too small to accommodate both attended and unattended stimuli.

Functional brain imaging studies in humans have extended these results by demonstrating increases in baseline activity not only at attended locations, but also in neural populations that respond preferentially to basic stimulus features such as motion (Chawla, Rees, and Friston 1999; Serences and Boynton 2007) and to more complex stimuli such as shapes (Stokes, Thompson, Nobre, and Duncan 2009) and objects (Peelen, Fei-Fei, and Kastner 2009). In these studies, baseline increases have been found at all stages of visual processing including the visual thalamus and striate cortex, which is somewhat inconsistent with reports from monkey physiology (which reported no baseline shifts in V1). One way to reconcile the apparent discrepancies between fMRI and physiology data is to consider the neural basis of the BOLD signal measured with fMRI. BOLD responses have been shown to more strongly correlate with local field potentials (LFPs) than with the action potentials that are typically recorded in physiology studies (Logothetis, Pauls, Augath, Trinath, and Oeltermann 2001; Logothetis and Wandell 2004). Critically, LFPs reflect multiple types of neuro-modulation including subthreshold changes in membrane potential, synaptic events, oscillatory activity, and after-potentials that follow action potentials. In addition, hemodynamic signals likely reflect the combined influence of small modulatory effects across large populations of neurons that may not be reliable at the level of single neuron recordings. Thus, BOLD fMRI might be more sensitive with respect to detecting the presence of an attentional modulation, at the expense of precise information about the origin of the modulatory signal.

The baseline shift and the early attentional modulation of stimulus-evoked responses represent complementary aspects of an early selection mechanism in monkey visual cortex. First, the baseline shift is consistent with a mechanism that increases the gain of the afferent volley of neural activity that is evoked by the presentation of a stimulus (Sylvester, Shulman, Jack, and Corbetta 2009). Thus, cells whose RFs correspond to an attended region of space already have a competitive advantage over cells whose RFs correspond to unattended regions of space, and they will tend to dominate the winner-take-all circuits posited by models such as biased competition (Desimone and Duncan 1995; Reynolds and Desimone 1999). Moreover, heightened spiking rates can be observed shortly after the presentation of an attended target, further amplifying the competitive advantage of relevant over irrelevant stimuli (e.g. Fig. 4.3a in Luck et al. 1997). Together, these findings are consistent with Broadbent's original conception that attention can use rudimentary qualities of the impinging sensory input such as location to influence attentional selection.

Taken at face value, the ERP, single cell, and fMRI studies reviewed in this section provide clear evidence for early signal gain in visual processing. Of course, these findings depend on linking propositions between observed neural activity and the ultimate behaviour of the organism that are often vague (Teller 1984). However, there is at least some evidence to support a correlation between enhanced sensory ERP components and decreased RTs in spatial cueing paradigms (e.g. Mangun et al. 1993; Zhang et al. 2012). While correlations cannot establish causal relationships, they substantially advance the argument that early modulation of neural activity directly contributes to behaviour. Moreover, caution must be exercised when interpreting the timing of the modulation effects. While it is widely accepted that the P1 and early modulations such as those documented by Luck et al. (1997) reflect feedforward processing, areas of parietal cortex (LIP) and frontal cortex can exert stimulus-driven responses with a latency on the order of 40–50 ms, raising the possibility that attentional feedback to early visual cortex could occur on an extremely short time scale (Bar et al. 2006; Bisley, Krishna, and Goldberg 2004). Nevertheless, it now seems clear that attentional modulation can occur well before the onset of neural markers that indicate more advanced neural processing, such as semantic analysis and the updating of information in working memory, processes that are typically thought to occur 300–400 ms post-stimulus (Kutas and Federmeier 2011; Kutas, Neville, and Holcomb 1987; Polich 2007).

Finally, it is important to point out that even though these studies are consistent with early selection accounts, they do not establish that early selection is the *only* mechanism of selective attention. Indeed, other ERP and single cell studies indicate that late selection can influence information processing as well. For example, early ERP components such as the P1 and N1 are similar for seen and unseen items in an attentional blink paradigm; however, large attenuations in later ERP components thought to index the updating of working memory are observed in the same context (Vogel, Luck, and Shapiro 1998). Thus, the emerging consensus from the cognitive neuroscience literature clarifies and solidifies earlier notions that both early and late selection can occur depending on task demands and are not mutually exclusive (Lavie and Tsal 1994; Vogel, Woodman, and Luck 2005; Yantis and Johnston 1990). However, the physiological evidence makes an especially compelling case indicating that pure inceptions of late selection theories can probably be entirely discarded.

Where and When of Noise Suppression via External Distractor Exclusion

Consider a classical visual search task, where subjects are presented with a display composed of green and red letters, e.g. T's and L's. The subject's task is to identify a target item composed of a conjunction of the two features present in the display, i.e. colour and shape (e.g. a red T) among the distractors (Treisman and Gelade 1980; see chapter by Wolfe). This type of task is designed to resemble real-world search situations, in which the visual system is constantly faced with an enormous amount of clutter from which behaviourally relevant information needs to be selected. Visual search entails at least two conceptually different processes: (i) the selection of the task-relevant stimulus and (ii) the filtering of multiple distractor stimuli that form the vast majority of information in the display. Thus far, we have focused almost exclusively on the anatomical and temporal properties of the target selection process. However, understanding the process of suppressing distractors is at least—if not more—important. Ultimately, these two processes are intimately linked, and there is now strong evidence that the neural mechanisms that support target selection and distractor exclusion interact to a high degree in visual cortex (e.g. Pinsk, Doniger, and Kastner 2004; Seidl, Peelen, and Kastner 2012). However, for reasons of clarity we will discuss the neural basis of distractor exclusion separately and without consideration of these interactions as they are not well understood at present.

The neural fate of unattended (distractor) information presents a second important question that arises from the early vs. late debate that we discussed above. Early selection accounts hold that unattended information should be mainly represented at early stages of processing, where the visual information is represented at a featural level, but not at later stages, where features are combined in complex ways to form objects and other semantically meaningful stimuli. In contrast, a late selection account would assume a neural representation at both early and advanced visual processing stages. Unfortunately, much less is known about the neural basis of distractor exclusion compared to the neural basis of target selection. This asymmetry is mainly due to the fact that most studies evaluate the effects of attention on a visual stimulus relative to a condition when the same stimulus is unattended. In this type of scenario, the effects of selective attention on unattended stimuli cannot be determined. Thus, some of the evidence discussed below will be indirect, but still important and telling about the mechanisms of distractor exclusion.

One example of such indirect evidence is the seminal work by Desimone and colleagues that led to the development of the biased competition model (Desimone and Duncan 1995; Kastner, De Weerd, Desimone, and Ungerleider 1998; Moran and Desimone 1985; Reynolds, Chelazzi, and Desimone 1999). As discussed above, the basic findings were two-fold. First, multiple stimuli appearing in a neuron's RF will interact in a mutually suppressive way suggesting that they are not processed independently. Instead, multiple stimuli engage in a competitive process that occurs automatically at the level of the RF. Second, when attention is allocated to one of two competing stimuli, the neural response to the attended stimulus is nearly as large as when the stimulus is presented in isolation (i.e. without the competing distractor). Thus, attention appears to restore the response of the attended stimulus and operates by counteracting the competing (suppressive) influence of the second stimulus. This finding suggests that attention not only boosts the response evoked by attended stimuli, but also acts to attenuate distractor interference by resolving competitive interactions between stimuli. Importantly, this mechanism of distractor exclusion has been shown to be tied to the RF and its immediate surround and scales with increasing RF size across the visual processing hierarchy (Kastner et al. 2001; Sundberg, Mitchell, and Reynolds 2009). However, as mentioned above, these studies did not directly measure the effects of attention on distractor stimuli, since only responses evoked by attended stimuli were considered. Thus, the conclusions drawn about the neural fate of distractor stimuli were based solely on changes in target-evoked responses as a function of different target–distractor configurations.

Recent studies using neuroimaging methods in human subjects have provided more direct evidence regarding the fate of unattended distractor stimuli. Some of the first evidence was obtained from studies in which the perceptual or cognitive load associated with a target stimulus was systematically and parametrically varied while responses evoked by the target and by distractors were separately assessed. In one study, the subject's attentional resources were parametrically modulated by manipulating the difficulty of a task at fixation while neural responses associated with an irrelevant and unattended peripheral motion stimulus were probed (Rees, Frith, and Lavie 1997) (see also chapter by Lavie). The response evoked by the unattended stimulus was significantly attenuated in motion-selective area MT when the task at fixation was difficult (high perceptual load) compared to when the task at fixation was easy (low perceptual load). O'Connor et al. (2002) extended these results by showing that responses evoked by an unattended visual stimulus also decreased in the thalamus, primary visual cortex, and extrastriate cortex depending on the load of a concurrent attentional task (see also Pinsk, et al. 2004; Schwartz et al. 2005). Interestingly, the load-dependent attenuation of distractor-related activity was strongest in the thalamus and extrastriate cortex and weaker in primary visual cortex, mirroring the pattern of effects observed for target-related attentional modulation (see section above, *The 'where' of selection via signal enhancement*). Thus, in addition to a gain control mechanism that modulates target-related neural signals throughout the visual processing hierarchy, active distractor suppression also appears to operate at multiple levels including the earliest stages of visual processing.

Corroborating evidence supporting an anatomically early locus of distractor suppression has also been found in studies of feature-based attention. One of the hallmarks of feature-based selection is that it operates by increasing neural responses to the selected feature (e.g. a particular direction of motion) regardless of spatial location, thus operating in a global manner across the visual field (Martinez-Trujillo and Treue 2004; Saenz, Buracas, and Boynton 2002; Serences and Boynton 2007; Treue and Martinez Trujillo 1999; see chapter by Scolari). In line with studies that examine the consequences of spatial selection on distractor processing, the response of neurons that are tuned to an attended feature (e.g. a direction of motion) are enhanced whereas the response of neurons tuned far from the attended feature are suppressed (Cohen and Maunsell 2011; Martinez-Trujillo and Treue 2004; Scolari, Byers, and Serences 2012; Serences, Saproo, Scolari, Ho, and Muftuler 2009). This joint enhancement and suppression occurs for all neurons that are sensitive to attended feature dimension, irrespective of whether the attended stimulus directly falls within their spatial receptive field. These results thus suggest that the exclusion of distractor information can happen at relatively early stages of sensory processing and that distractor suppression is not strictly tied to acts of selection based on spatial location.

In general, the spatial and feature-based suppression of distractor-related activity is more compatible with early selection accounts as it occurs at early stages of the visual hierarchy (although the timing of these effects is largely an open question). However, these studies all dealt with basic visual features as opposed to more complex objects and experimental settings that more closely resemble real-world search scenarios. Interestingly, and in line with studies that use basic visual features, most studies that examine the influence of attention on the processing of more complex objects reveal little representation of unattended distractor categories. In one classic study, a house and a face stimulus were presented simultaneously in separate hemifields and subjects were cued to attend either to the location of the face or to the location of the house (Serences, Schwarzbach, Courtney, Golay, and Yantis 2004; Wojciulik, Kanwisher, and Driver 1998). When subjects attended to the face stimuli, face-selective regions of ventral visual cortex were more responsive than house-selective regions of ventral visual cortex. The opposite was true when subjects attended to houses. More recent studies have extended these findings to more naturalistic and thus complex scenarios by showing that neural activity in object-selective cortex is entirely dominated by task-related demands when subjects extract categorical information from natural scenes. In one study, subjects attended to briefly presented street scenes and detected the presence of people or cars in the scenes. In object-selective ventral visual cortex, only task-relevant information was processed to the categorical level even when the relevant information was not spatially attended. In contrast, task-irrelevant information was not processed to the categorical level even when it was spatially attended (Peelen et al. 2009). Subsequent studies demonstrated that task-irrelevant objects can sometimes be processed up to the categorical level; however, their representation is weaker than that of task-relevant objects (Seidl et al. 2012). In addition, distractors that were task-relevant in an immediately preceding trial that become task-irrelevant in the present trial are also actively suppressed (Fig. 4.2). Together, these and other related findings strongly suggest that distractor information is largely filtered out at the level of object-selective cortex, thereby further supporting early selection accounts of attentional modulation.

While the pattern of results regarding the fate of distractors is largely consistent across single-unit and neuroimaging studies, results from patient literature are more mixed with respect to early and late selection accounts. For example, the BOLD response evoked by faces and by common objects in patients suffering from visuo-spatial hemineglect was similar when the stimuli were presented to the neglected hemifield and when the stimuli were presented to the intact hemifield (Rees et al. 2000; Vuilleumier et al. 2001). The similarity in magnitude of the response suggests that even unattended

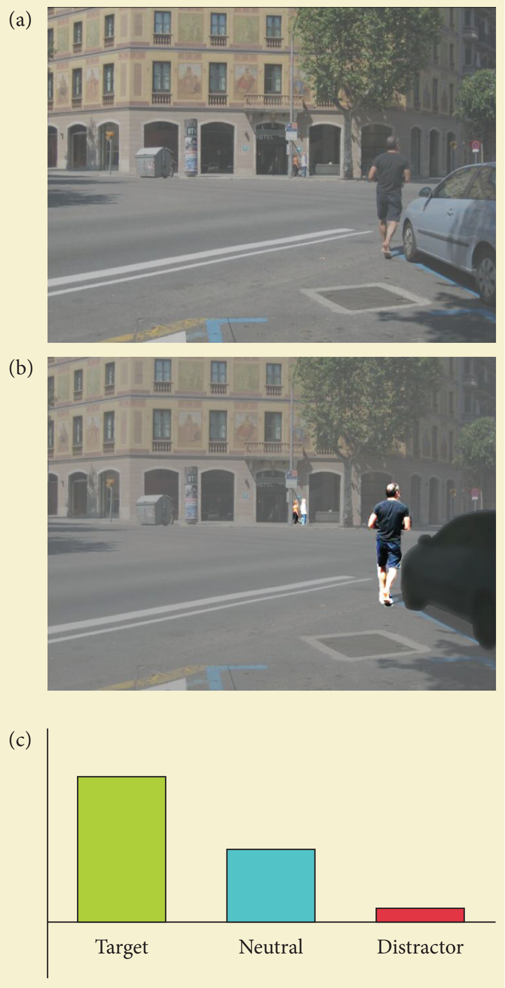

FIGURE 4.2 Object category-based selection from natural scenes. In the scene (a), people are the category that is relevant to ongoing behaviour (i.e. target category), and cars are the object category that was previously relevant but is presently not relevant (i.e. distractor category), whereas all other object categories present in the scene are never task-relevant (i.e. neutral category, such as trees or houses). Visual search in natural scenes is accomplished through a combination of enhancing task-relevant information and suppressing a previous attentional set relative to processing of neutral categories. The resulting representation of object categories in object-selective cortex is schematically shown in the modified scene (b) relative to the original scene (a). (A) Adapted from Seidl, K. N., Peelen, M. V., and Kastner, S., Neural evidence for distractor suppression during visual search in real-world scenes, *Journal of Neuroscience*, 32, pp. 11812–11819 © 2012, The Society for Neuroscience.

stimuli undergo processing to a relatively advanced stage of processing, perhaps even to the level of category-specific object representations. However, in light of previous findings suggesting that object-based selection operates in a spatially global manner (similar to feature-based attention, e.g. Peelen et al. 2009), it is possible that objects presented in the intact hemifield heavily influenced the representation in the lesioned hemisphere. On this account, object identity was never represented in the lesioned hemisphere. Instead, activation observed in the lesioned hemisphere may have been related to the global processing of the attended target. Other lesion work suggests that neglected—and thus unattended—stimuli undergo significant processing, at least to the point at which they can influence behaviour (Marshall and Halligan 1988). For instance, neglect patients show behavioural interference effects tied to the presentation of incompatible distractors in their neglected hemifield, even though they report little awareness of the competing distractors (Shomstein, Kimchi, Hammer, and Behrmann 2010). Thus, it appears that at least some information about stimuli from the neglected hemifield can be processed; however, the depth of processing is not entirely clear, nor is the influence of parametric manipulations of task difficulty. Future studies will be required to determine the extent to which these late selection effects are still observed when sufficient attentional resources are allocated to the non-neglected side of space.

While relatively little is known about the fate of unattended distractor stimuli, even less is known about the temporal dynamics of distractor exclusion (or the 'when' of distractor exclusion). One possibility is that the processing of target-related information is simply sped up relative to the processing of distractor information. Such changes in the efficiency of processing may be reflected in the latency at which the target and distractor-evoked responses are registered in the visual system. To evaluate this possibility, several single-unit recording studies examined the response latency associated with identical stimuli when they were attended or unattended. With the exception of one study, the results suggest that attention does not significantly alter the response latency of single neurons (Bisley et al. 2004; Cook and Maunsell 2004; Lee, Williford, and Maunsell 2007; Reynolds, Pasternak, and Desimone 2000). In the one study that did find a positive effect, the latency shift was found both at the level of single neurons and at the level of local populations (Sundberg, Mitchell, Gawne, and Reynolds 2012), and was on the order of 1–2 msec. Despite this relatively modest shift, the decreased response latency may still have a significant impact on the computations performed by a network or neurons given that much of the information in a neural code likely depends on spike timing. Furthermore, small differences measured at the single cell or local population level can easily translate into more robust differences when averaged across larger neuronal populations, such as those observed using scalp EEG or electrocorticography (ECoG). Interestingly, a recent ERP study demonstrated that the enhancement of target-related activity measured over visual cortex started 220 msec after the onset of an attentional cue, whereas the suppression of distractor-evoked responses lagged by an additional 130 msec (Andersen and Muller 2010). Thus, while the temporal dynamics of distractor suppression are still not well understood, both of these studies highlight the importance of addressing the interaction between the facilitatory and suppressive effects of selective attention.

WHERE AND WHEN OF ATTENTION-RELATED INTERNAL NOISE REDUCTION

Signal enhancement can be used to efficiently encode relevant stimuli and to attenuate the influence of external distractors on information processing. In addition, recent research—primarily using single-unit physiology—also demonstrates that selective attention can reduce the intrinsic internal noise observed in spiking activity as quantified by a reduction in the ratio of the mean response to the variance of the response (termed the *Fano Factor*; Mitchell, Sundberg, and Reynolds 2007). Such internal noise reduction likely plays a complementary role to signal enhancement, as the signal-to-noise ratio (SNR) of neurons increases with gain and also with a reduction in noise (Fig. 4.3). This may be particularly important in light of neurobiologically observed

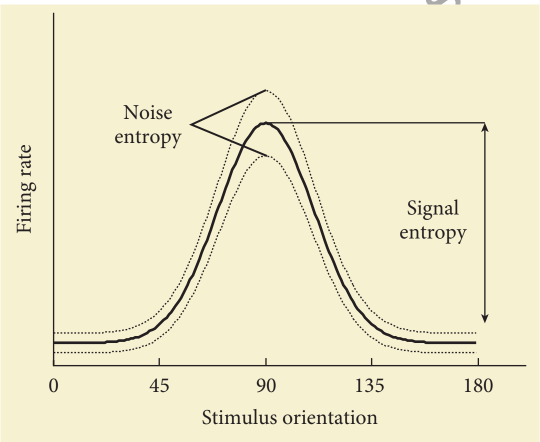

FIGURE 4.3 The amount of information that a sensory neuron encodes about sensory features depends on both response gain and response variance. Depicted here is a cartoon schematic of the tuning function of a simple cell in primary visual cortex that responds maximally to a 90° oriented bar of light. Mutual Information (or MI; Shannon 1949) is a useful metric for quantifying how much information about a sensory stimulus can be gained by measuring the response of this neuron to a visual stimulus. MI is, intuitively speaking, the difference between the signal entropy—or the amount of variance in the neural response that is systematically related to changes in the stimulus—and the noise entropy—or the amount of variance in the neural response that is unrelated to changes in the stimulus (i.e. 'noise'). Increasing the gain of sensory neurons will increase the signal entropy by increasing the range over which responses vary as a function of the stimulus. All else being equal, this will increase the MI between neural responses and the stimulus features being encoded. Alternatively, decreasing the variability of a neuron (i.e. decreasing the Fano Factor) will reduce the amount of variability in neural responses that is unrelated to changes in the stimulus. All else being equal, this will also increase the MI between neural responses and the stimulus features being encoded.

neuronal noise, which often far exceeds the typical assumption of a Poisson distribution (where the variance is equivalent to the mean; Anderson, Mitchell, and Reynolds 2011). In extreme cases where the Poisson expectation is far exceeded, increasing gain will have little effect on the information conveyed by a single unit about an attended stimulus feature because the noise will increase at a fast enough rate to offset any gain-related improvements in SNR. In these situations, reducing the variability of single unit responses may be a critical component to ensuring that relevant sensory stimuli get processed more efficiently than irrelevant distractors.

Even though attention can attenuate trial-by-trial variability in single-unit responses, a considerable amount of noise remains in the system. Thus, some pooling or averaging of responses is required to form stable representations of relevant sensory features. However, averaging cannot remove noise that is correlated between similarly tuned neurons, thus placing a limit on the precision of sensory representations that are based on population codes (Abbott and Dayan 1999; Averbeck, Latham, and Pouget 2006; Averbeck and Lee 2006; Bair, Zohary, and Newsome 2001; Johnson 1980; Kohn and Smith 2005; Mitchell et al. 2007; Shadlen, Britten, Newsome, and Movshon 1996; Shadlen and Newsome 1994). Given the potential limits imposed by correlated noise, several recent studies have focused on uncovering the role of attention in decreasing shared noise between similarly tuned neurons to improve the efficacy of pooling sensory signals. For example, Mitchell and co-workers (Mitchell, Sundberg, and Reynolds 2009) demonstrated that spatial attention can attenuate correlated noise, particularly low-frequency noise, in mid-level area V4. Assuming a simple population read-out rule that is based on averaging responses across all neurons tuned to the attended stimulus, the observed decorrelations can lead to an estimated 40% improvement in the SNR of the neuronal population, as compared to an estimated 10% improvement in SNR due solely to increases in gain (Mitchell et al. 2009). While the exact benefit in terms of SNR needs further exploration under different task conditions and assumptions about how population codes are used to information perceptual decisions, this asymmetry highlights the powerful role that changes in variability and covariance can play in influencing the quality of sensory representations. Further underscoring the importance of noise correlations, Cohen and Maunsell (Cohen and Maunsell 2009, 2010, 2011) demonstrated decorrelations related to both space- and feature-based attention in areas V4 and MT. Importantly, the degree of decorrelation predicted trial-by-trial fluctuations in behavioural performance.

While this field of inquiry is nascent, and attention-mediated changes in variability and covariance have only been documented in a few mid-level structures, even less is known about the time course of these modulations. In one relevant experiment, Mitchell et al. (2007, 2009) used a stimulus that moved into the receptive field of a V4 neuron for 1000 ms before leaving the receptive field. Thus, the authors could examine changes in the Fano Factor and in covariance across counting windows that varied in size from approximately 10 msec to 400 msec. They observed that both the attention-related reduction in the Fano Factor and the attention-related reduction in covariance onset quite rapidly. In fact, spatial attention seemed to anticipate the trajectory of the moving stimulus, and significant decreases in internal noise were already evident at the point that the stimulus entered a neuron's RF. This onset latency is early enough in time to influence decision-making in speeded perceptual tasks, consistent with the documented relationship between covariability and accuracy on a cueing task (Cohen and Maunsell 2009, 2011). In addition, the magnitude of attention-related reductions in variance and covariance increased as the size of the counting window increased, suggesting that attention had the largest modulatory impact on low-frequency noise.

Collectively, these studies thus suggest that attention improves the efficiency of information processing in large part via changes in the variance and covariance of neurons in mid-level visual areas. Moreover, the temporally early onset of these effects suggests that changes in the noise characteristics of sensory neurons might operate relatively early in information processing during the extraction of basic low-level features. However, given that only a handful of studies have been done in this area there is still much to learn about the anatomical extent of these noise modulations as well as their onset time and the temporal window over which they operate.

LATE SELECTION VIA THE EFFICIENT READ-OUT OF SENSORY SIGNALS

Although most of the discussion surrounding early versus late selection has centred on the anatomical locus, timing, and nature of sensory gain modulations, a parallel line of reasoning proposes that many of the effects attributed to selective attention can be explained without invoking either early or late modulations in early sensory regions. Instead, these *selective read-out* models posit that decision mechanisms can pool responses selectively from neurons that are optimally tuned to discriminate the attended stimulus, and that this selective pooling is sufficient to improve information processing even in the complete absence of sensory gain modulations (Dosher, Liu, Blair, and Lu 2004; Eckstein, Peterson, Pham, and Droll 2009; Eckstein, Thomas, Palmer, and Shimozaki 2000; Palmer, Verghese, and Pavel 2000; Shaw 1984). While it is clear that sensory gain modulations do occur as the result of attention, these models nevertheless make a critical, and often overlooked, point about the importance of maximizing the efficiency of how populations of sensory neurons are read out during decision-making (Fig. 4.4).

The power of selective read-out is perhaps most evident when considering how attention can greatly attenuate—or sometimes even eliminate—the influence of irrelevant distracting items that compete with relevant stimuli (Palmer and Moore 2009; Yigit-Elliott, Palmer, and Moore 2011). Since these models posit that decision mechanisms only read out responses from sensory neurons that are optimally tuned to make the relevant discrimination, the influence of irrelevant distracting items is automatically attenuated as responses associated with these stimuli have no impact whatsoever on the decision process. In this manner, selective read-out mechanisms can efficiently shunt

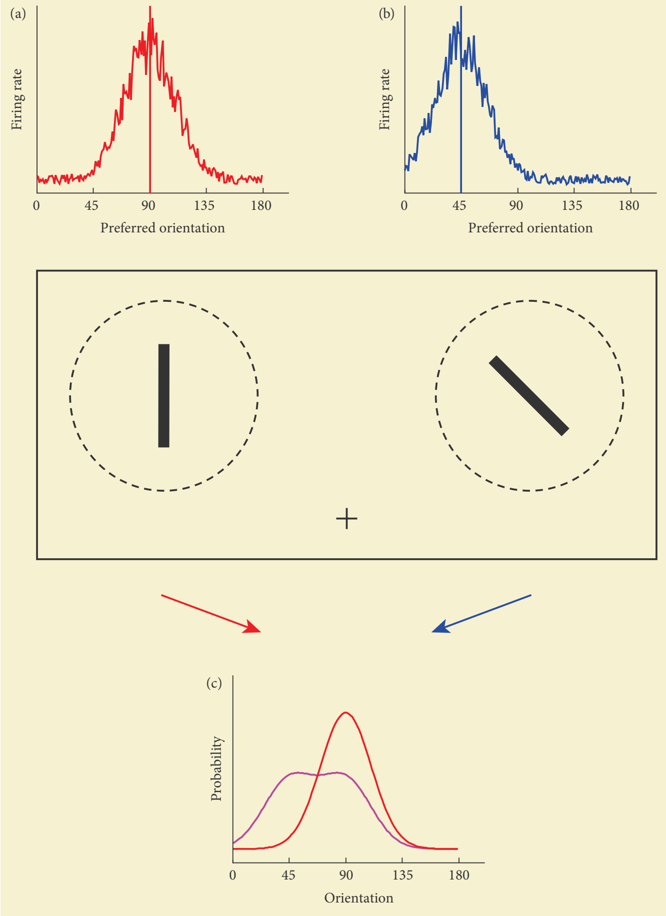

FIGURE 4.4 Selectively basing perceptual decisions on neural populations that are most sensitive to the relevant stimulus can improve the speed and accuracy of behavioural responses. (bottom panel) Schematic of a stimulus display presented briefly to a subject. The subject's task is to report, as quickly and accurately as possible, whether or not a vertical line is present in the display. (a) The response profile evoked by the vertical 90° stimulus across a population of sensory neurons that are maximally responsive (i.e. 'tuned to') different orientations and that have a spatial receptive field in the left hemifield. (b) The response profile evoked by a 45° stimulus across a different population of sensory neurons that have a spatial receptive field in the right hemifield. (c—Magenta line) If the subject has no advance knowledge about where the vertical line might appear, then they might reasonably pool responses across both neural populations before assessing the likelihood that a vertical line was present. This will lead to a relatively blurred and imprecise estimate of the stimulus features that gave rise to the sensory responses profiles show in panels a and b. (c—Red line) In contrast, if the subject knows in advance that the vertical line, if present, will appear in the left hemifield, then they need only base their likelihood estimate on the output of the neurons that have a spatial receptive field in the left hemifield (the neural response profile depicted in panel (a)). A comparison of the magenta and red lines suggests that inferences in this case will be more precise because irrelevant information provided by the neurons shown in panel b is not allowed to influence decision-making. Thus, selectively basing decisions on the most sensitive neurons can improve the efficiency of information processing, even in the absence of any direct modulation of sensory responses (e.g. gain or variance modulation).

interference from sensory neurons that encode irrelevant features, thereby increasing the discriminability of neural signals associated with attended stimuli.

While the selective read-out framework is appealing because complex modulations of sensory responses are not required, most of the work in this field is based on pure theory and mathematical models. Thus, even though selective read-out models can account for a variety of observed attentional modulations without recourse to sensory gain, this usually amounts to a proof of concept as opposed to a proof of existence. However, one recent neuroimaging study provides direct evidence that selective read-out plays an important role in implementing efficient information processing. Pestilli and co-workers (Pestilli, Carrasco, Heeger, and Gardner 2011) had observers search for a target embedded in a set of distractors, and demonstrated that sensory gain alone was not sufficient to account for the observed improvement in behaviour with attention. Instead, the data suggest that decision mechanisms over-weighted sensory responses associated with the attended stimulus and under-weighted responses associated with unattended distractors. This biased read-out process thereby ensured that signals associated with the attended stimulus dominated the decision process whereas signals associated with distractors had little or no impact. Again, the key element of this account is that selection occurs when sensory signals are integrated and evaluated by late-stage decision mechanisms, long after basic stimulus features are fully analyzed.

Related work has been done in the context of perceptual learning, where improvements in behaviour with practice are often thought to involve changes in the optimality of sensory read-out (Law and Gold 2008, 2009; Petrov, Dosher, and Lu 2005). For instance, Law and Gold (2008) found that firing rates in MT neurons did not significantly change after extensive training on a motion discrimination task. However, there were concurrent changes in the firing rates of the neurons in the lateral intraparietal area (LIP), an area implicated in accumulating sensory evidence during decision-making (Law and Gold 2008; see also Gold and Shadlen 2007). These data suggest that read-out from informative sensory neurons plays a more important role—at least in their task—than changes in sensory gain. In addition, even though the study ostensibly examined perceptual learning, Law and Gold (2008) used a motion stimulus that moved in a direction that was tailored to the tuning preference of the MT cells that they isolated each day. Thus, their task deviates from classic studies of learning that use a fixed stimulus feature, and may more closely resemble a more general perceptual task used in many attention studies (e.g. Treue and Maunsell 1996). The results are also consistent with some models of selective attention (Palmer and Moore 2009; Palmer et al. 2000), in which the most sensitive MT responses are pooled with linear weights and uninformative neurons are filtered out, leading to lower discrimination thresholds (Gold, Law, Connolly, and Bennur 2010; Law and Gold 2009).

While these empirical studies and the associated theories provide evidence for the importance of biased read-out as a mechanism of late selection, many important questions remain. For example, it seems likely that the importance of selective read-out depends critically on the number of competing items in the display. If only one item is present, then there is no need to selectively read out signals from only the attended item, so increasing the gain of sensory responses should be adequate to achieve whatever signal-to-noise ratio is required to perform the required perceptual task. However, the benefits of sensory gain will likely decrease as the number of competing distractors increases, at least under the reasonable assumption that gain cannot be increased indefinitely. In this situation, shunting the influence of distractors by discounting their impact on a late decision mechanism is a computationally and metabolically efficient way to implement selection and to ensure that only the most relevant sensory stimuli influence working memory and subsequent behavioural responses. Selective read-out therefore provides an example of a robust late-selection mechanism that, especially when combined with early sensory gain, highlights the increasingly accepted idea that selection is achieved through a combination of mechanisms that simultaneously operate at nearly all stages of information processing.

CONCLUSIONS

Over the last several decades, data from neuroimaging and electrophysiology studies have shown that the neural mechanisms of selective attention operate at almost all stages of the visual system, as well as in many areas outside of classically defined visual cortex. When viewed together, these studies firmly establish that selection is neither early nor late. Instead, the locus of selection, both in terms of anatomy and time, flexibly depends on the demands placed on sensory processing machinery by the behavioural goals of an observer. Tasks that require highly focused attention on a specific location or feature will encourage early selection, whereas less demanding tasks that can be performed with a more diffuse attentional focus will accommodate late selection. Finally, a complete understanding of selective information processing is not all about understanding where and when the gain of sensory neurons is modulated: emerging evidence suggests equally important roles for modulating the variance and covariance of sensory neurons, and for selectively reading out information from only the most informative sensory neurons.

Moving ahead, the major challenge for future investigators is to meld the multiple mechanisms that support selective attention into a unified framework. For instance, it is well accepted that each stage of the visual processing hierarchy contributes differently to visual perception. Therefore, it also seems likely that selective attention implements different functions in different visual areas, and that each function is determined by the specific processing capabilities of a region. For example, at the thalamic level, attention may serve to control neural response gain. At early cortical processing stages, attention may influence feature and context selective responses, which may facilitate the basic mechanisms of scene segmentation and grouping (e.g. Ito and Gilbert 1999). At intermediate cortical processing stages where multiple stimuli often fall within a single RF, attention may mediate the filtering of unwanted information through internal and external noise reduction (i.e. distractor exclusion). These diverse modulatory processes appear to be controlled by a higher-order fronto-parietal network of brain areas that may coordinate large-scale attentional modulation via subcortical structures such as the superior colliculus (Zenon and Krauzlis 2012; see chapter by Krauzlis) and the pulvinar nucleus of the thalamus (Saalmann, Pinsk, Wang, Li, and Kastner 2012; see chapter by Saalmann). Ultimately, however, these brain systems cooperate to select relevant visual information for further processing in memory and other cognitive systems in order to guide actions. In this respect, attention can be described as a multi-level selection process that operates by integrating information across multiple anatomical and temporal scales to achieve behavioural goals in a flexible and adaptive manner.

ACKNOWLEDGEMENTS

Supported by NIH R01-MH092345 and a James S. McDonnell Foundation Scholar Award to J.T.S. and by NIH RO1-MH64043, RO1-EY017699, R21-EY0211078, and NSF BCS-1025149 to S.K.

REFERENCES

Abbott, L. F. and Dayan, P. (1999). The effect of correlated variability on the accuracy of a population code. *Neural Computation 11*: 91–101.

Andersen, S. K. and Muller, M. M. (2010). Behavioral performance follows the time course of neural facilitation and suppression during cued shifts of feature-selective attention. *Proceedings of the National Academy of Sciences USA 107*: 13878–13882.

Anderson, E. B., Mitchell, J. F., and Reynolds, J. H. (2011). Attentional modulation of firing rate varies with burstiness across putative pyramidal neurons in macaque visual area V4. *Journal of Neuroscience 31*: 10983–10992.

Anton-Erxleben, K., Abrams, J., and Carrasco, M. (2011). Equality judgments cannot distinguish between attention effects on appearance and criterion: A reply to Schneider (2011). *Journal of Vision 11*. doi: 10.1167/11.13.8.

Averbeck, B. B., Latham, P. E., and Pouget, A. (2006). Neural correlations, population coding and computation. *Nature Reviews Neuroscience 7*: 358–366.

Averbeck, B. B. and Lee, D. (2006). Effects of noise correlations on information encoding and decoding. *Journal of Neurophysiology 95*: 3633–3644.

Bair, W., Zohary, E., and Newsome, W. T. (2001). Correlated firing in macaque visual area MT: Time scales and relationship to behavior. *Journal of Neuroscience 21*: 1676–1697.

Bar, M., Kassam, K. S., Ghuman, A. S., Boshyan, J., Schmid, A. M., Dale, A. M., Hamalainen, M. S., Marinkovic, K., Schacter, D. L., Rosen, B. R., and Halgren, E. (2006). Top-down facilitation of visual recognition. *Proceedings of the National Academy of Sciences USA 103*: 449–454.

Bisley, J. W., Krishna, B. S., and Goldberg, M. E. (2004). A rapid and precise on-response in posterior parietal cortex. *Journal of Neuroscience 24*: 1833–1838.

Broadbent, D. (1958). *Perception and Communication*. London: Pergamon.

Buffalo, E. A., Fries, P., Landman, R., Liang, H., and Desimone, R. (2010). A backward progression of attentional effects in the ventral stream. *Proceedings of the National Academy of Sciences USA 107*: 361–365.

Carrasco, M. (2009). Cross-modal attention enhances perceived contrast. *Proceedings of the National Academy of Sciences USA 106*: 22039–22040.

Carrasco, M., Fuller, S., and Ling, S. (2008). Transient attention does increase perceived contrast of suprathreshold stimuli: A reply to Prinzmetal, Long, and Leonhardt (2008). *Perception & Psychophysics 70*: 1151–1164.

Carrasco, M., Ling, S., and Read, S. (2004). Attention alters appearance. *Nature Neuroscience 7*: 308–313.

Carrasco, M. and McElree, B. (2001). Covert attention accelerates the rate of visual information processing. *Proceedings of the National Academy of Sciences USA 98*: 5363–5367.

Chawla, D., Rees, G., and Friston, K. J. (1999). The physiological basis of attentional modulation in extrastriate visual areas. *Nature Neuroscience 2*: 671–676.

Cherry, E. C. (1953). Some experiments on the recognition of speech, with one and with two ears. *Journal of Acoustic Society of America 25*: 975–979.

Cohen, M. R. and Maunsell, J. H. (2009). Attention improves performance primarily by reducing interneuronal correlations. *Nature Neuroscience 12*: 1594–1600.

Cohen, M. R. and Maunsell, J. H. (2010). A neuronal population measure of attention predicts behavioral performance on individual trials. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 30*: 15241–15253.

Cohen, M. R. and Maunsell, J. H. (2011). Using neuronal populations to study the mechanisms underlying spatial and feature attention. *Neuron 70*: 1192–1204.

Connor, C. E., Preddie, D. C., Gallant, J. L., and Van Essen, D. C. (1997). Spatial attention effects in macaque area V4. *Journal of Neuroscience 17*: 3201–3214.

Cook, E. P. and Maunsell, J. H. (2002a). Attentional modulation of behavioral performance and neuronal responses in middle temporal and ventral intraparietal areas of macaque monkey. *Journal of Neuroscience 22*: 1994–2004.

Cook, E. P. and Maunsell, J. H. (2002b). Dynamics of neuronal responses in macaque MT and VIP during motion detection. *Nature Neuroscience 5*: 985–994.

Cook, E. P. and Maunsell, J. H. (2004). Attentional modulation of motion integration of individual neurons in the middle temporal visual area. *Journal of Neuroscience 24*: 7964–7977.

Corteen, R. S. and Wook, B. (1972). Autonomic responses to shock-associated words in an unattended channel. *Journal of Experimental Psychology 94*: 308–313.

Desimone, R. and Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience 18*: 193–222.

Deutsch, J. and Deutsch, D. (1963). Attention: Some theoretical considerations. *Psychological Review 70*: 80–90.

Dosher, B. A., Liu, S. H., Blair, N., and Lu, Z. L. (2004). The spatial window of the perceptual template and endogenous attention. *Vision Research 44*: 1257–1271.

Duncan, J. (1980). The locus of interference in the perception of simultaneous stimuli. *Psychological Review 87*: 272–300.

Eckstein, M. P., Peterson, M. F., Pham, B. T., and Droll, J. A. (2009). Statistical decision theory to relate neurons to behavior in the study of covert visual attention. *Vision Research 49*: 1097–1128.

Eckstein, M. P., Thomas, J. P., Palmer, J., and Shimozaki, S. S. (2000). A signal detection model predicts the effects of set size on visual search accuracy for feature, conjunction, triple conjunction, and disjunction displays. *Perception & Psychophysics 62*: 425–451.

Eriksen, B. A. and Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception & Psychophysics 16*: 143–149.

Francolini, C. M. and Egeth, H. E. (1980). On the nonautomaticity of 'automatic' activation: Evidence of selective seeing. *Perception & Psychophysics 27*: 331–342.

Fuller, S., Park, Y., and Carrasco, M. (2009). Cue contrast modulates the effects of exogenous attention on appearance. *Vision Research 49*: 1825–1837.

Gandhi, S. P., Heeger, D. J., and Boynton, G. M. (1999). Spatial attention affects brain activity in human primary visual cortex. *Proceedings of the National Academy of Sciences USA 96*: 3314–3319.

Gold, J. I., Law, C. T., Connolly, P., and Bennur, S. (2010). Relationships between the threshold and slope of psychometric and neurometric functions during perceptual learning: Implications for neuronal pooling. *Journal of Neurophysiology 103*: 140–154.

Gold, J. I. and Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience 30*: 535–574.

Gottlieb, J. P., Kusunoki, M., and Goldberg, M. E. (1998). The representation of visual salience in monkey parietal cortex. *Nature 391*: 481–484.

Haenny, P. E. and Schiller, P. H. (1988). State dependent activity in monkey visual cortex. I: Single cell activity in V1 and V4 on visual tasks. *Experimental Brain Research/Experimentelle Hirnforschung/Experimentation Cerebrale 69*: 225–244.

Hillyard, S. A. and Anllo-Vento, L. (1998). Event-related brain potentials in the study of visual selective attention. *Proceedings of the National Academy of Sciences USA 95*: 781–787.

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. *Philosophical Transactions of the Royal Society of London B: Biological Sciences 353*: 1257–1270.

Ito, M. and Gilbert, C. D. (1999). Attention modulates contextual influences in the primary visual cortex of alert monkeys. *Neuron 22*: 593–604.

Johannes, S., Munte, T. F., Heinze, H. J., and Mangun, G. R. (1995). Luminance and spatial attention effects on early visual processing. *Brain Research: Cognitive Brain Research 2*: 189–205.

Johnson, K. O. (1980). Sensory discrimination: Decision process. *Journal of Neurophysiology 43*: 1771–1792.

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.

Kastner, S., De Weerd, P., Pinsk, M. A., Elizondo, M. I., Desimone, R., and Ungerleider, L. G. (2001). Modulation of sensory suppression: Implications for receptive field sizes in the human visual cortex. *Journal of Neurophysiology 86*: 1398–1411.

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

Kelly, S. P., Gomez-Ramirez, M., and Foxe, J. J. (2008). Spatial attention modulates initial afferent activity in human primary visual cortex. *Cerebral Cortex 18*: 2629–2636.

Khoe, W., Mitchell, J. F., Reynolds, J. H., and Hillyard, S. A. (2005). Exogenous attentional selection of transparent superimposed surfaces modulates early event-related potentials. *Vision Research 45*: 3004–3014.

Kohn, A. and Smith, M. A. (2005). Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *Journal of Neuroscience 25*: 3661–3673.

Kutas, M. and Federmeier, K. D. (2011). Thirty years and counting: Finding meaning in the N400 component of the event-related brain potential (ERP). *Annual Review of Psychology 62*: 621–647.

Kutas, M., Neville, H. J., and Holcomb, P. J. (1987). A preliminary comparison of the N400 response to semantic anomalies during reading, listening and signing. *Electroencephalography and Clinical Neurophysiology. Supplement 39*: 325–330.

Lavie, N. and Tsal, Y. (1994). Perceptual load as a major determinant of the locus of selection in visual attention. *Perception & Psychophysics 56*: 183–197.

Law, C. T. and Gold, J. I. (2008). Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. *Nature Neuroscience 11*: 505–513.

Law, C. T. and Gold, J. I. (2009). Reinforcement learning can account for associative and perceptual learning on a visual-decision task. *Nature Neuroscience 12*: 655–663.

Lee, J., Williford, T., and Maunsell, J. H. (2007). Spatial attention and the latency of neuronal responses in macaque area V4. *Journal of Neuroscience 27*: 9632–9637.

Liu, T., Abrams, J., and Carrasco, M. (2009). Voluntary attention enhances contrast appearance. *Psychological Science: A Journal of the American Psychological Society/APS 20*: 354–362.

Liu, T., Fuller, S., and Carrasco, M. (2006). Attention alters the appearance of motion coherence. *Psychonomic Bulletin & Review 13*: 1091–1096.

Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., and Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature 412*: 150–157.

Logothetis, N. K. and Wandell, B. A. (2004). Interpreting the BOLD signal. *Annual Review of Physiology 66*: 735–769.

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

McAdams, C. J. and Maunsell, J. H. (1999). Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *Journal of Neuroscience 19*: 431–441.

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* (vol. 14, pp. 219–243). Cambridge, Mass.: MIT Press.

Marshall, J. C. and Halligan, P. W. (1988). Blindsight and insight in visuo-spatial neglect. *Nature 336*: 766–767.

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 Neuroscience 2*: 364–369.

Martinez-Trujillo, J. C. and Treue, S. (2004). Feature-based attention increases the selectivity of population responses in primate visual cortex. *Current Biology: CB 14*: 744–751.

Mitchell, J. F., Sundberg, K. A., and Reynolds, J. H. (2007). Differential attention-dependent response modulation across cell classes in macaque visual area V4. *Neuron 55*: 131–141.

Mitchell, J. F., Sundberg, K. A., and Reynolds, J. H. (2009). Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron 63*: 879–888.

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

Moray, N. (1959). Attention in dichotic listening: Affective cues and the influence of instructions. *Quarterly Journal of Experimental Psychology 11*: 56–60.

Motter, B. C. (1993). Focal attention produces spatially selective processing in visual cortical areas V1, V2, and V4 in the presence of competing stimuli. *Journal of Neurophysiology 70*: 909–919.

Neisser, U. and Becklen, R. (1975). Selective looking: Attending to visually specified events. *Cognitive Psychology 7*: 480–494.

Noesselt, T., Hillyard, S. A., Woldorff, M. G., Schoenfeld, A., Hagner, T., Jancke, L., Tempelmann, C., Hinrichs, H., and Heinze, H. J. (2002). Delayed striate cortical activation during spatial attention. *Neuron 35*: 575–587.

O'Connor, D. H., Fukui, M. M., Pinsk, M. A., and Kastner, S. (2002). Attention modulates responses in the human lateral geniculate nucleus. *Nature Neuroscience 5*: 1203–1209.

Palmer, J. and Moore, C. M. (2009). Using a filtering task to measure the spatial extent of selective attention. *Vision Research 49*: 1045–1064.

Palmer, J., Verghese, P., and Pavel, M. (2000). The psychophysics of visual search. *Vision Research 40*: 1227–1268.

Pashler, H. (1994). Dual-task interference in simple tasks: Data and theory. *Psychological Bulletin 116*: 220–244.

Peelen, M. V., Fei-Fei, L., and Kastner, S. (2009). Neural mechanisms of rapid natural scene categorization in human visual cortex. *Nature 460*: 94–97.

Pestilli, F., Carrasco, M., Heeger, D. J., and Gardner, J. L. (2011). Attentional enhancement via selection and pooling of early sensory responses in human visual cortex. *Neuron 72*: 832–846.

Petrov, A. A., Dosher, B. A., and Lu, Z. L. (2005). The dynamics of perceptual learning: An incremental reweighting model. *Psychological Review 112*: 715–743.

Pinsk, M. A., Doniger, G. M., and Kastner, S. (2004). Push–pull mechanism of selective attention in human extrastriate cortex. *Journal of Neurophysiology 92*: 622–629.

Polich, J. (2007). Updating P300: An integrative theory of P3a and P3b. *Clinical Neurophysiology 118*: 2128–2148.

Pouget, A., Dayan, P., and Zemel, R. S. (2003). Inference and computation with population codes. *Annual Review of Neuroscience 26*: 381–410.

Prinzmetal, W., Long, V., and Leonhardt, J. (2008). Involuntary attention and brightness contrast. *Perception & Psychophysics 70*: 1139–1150.

Rees, G., Frith, C. D., and Lavie, N. (1997). Modulating irrelevant motion perception by varying attentional load in an unrelated task. *Science 278*: 1616–1619.

Rees, G., Wojciulik, E., Clarke, K., Husain, M., Frith, C., and Driver, J. (2000). Unconscious activation of visual cortex in the damaged right hemisphere of a parietal patient with extinction. *Brain 123*(Pt 8): 1624–1633.

Ress, D., Backus, B. T., and Heeger, D. J. (2000). Activity in primary visual cortex predicts performance in a visual detection task. *Nature Neuroscience 3*: 940–945.

Reynolds, J. H., Chelazzi, L., and Desimone, R. (1999). Competitive mechanisms subserve attention in macaque areas V2 and V4. *Journal of Neuroscience 19*: 1736–1753.

Reynolds, J. H. and Desimone, R. (1999). The role of neural mechanisms of attention in solving the binding problem. *Neuron 24*: 19–29, 111–125.

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

Roelfsema, P. R., Lamme, V. A., and Spekreijse, H. (1998). Object-based attention in the primary visual cortex of the macaque monkey. *Nature 395*: 376–381.

Saalmann, Y. B., Pinsk, M. A., Wang, L., Li, X., and Kastner, S. (2012). The pulvinar regulates information transmission between cortical areas based on attention demands. *Science 337*: 753–756.

Saenz, M., Buracas, G. T., and Boynton, G. M. (2002). Global effects of feature-based attention in human visual cortex. *Nature Neuroscience 5*: 631–632.

Schneider, K. A. (2006). Does attention alter appearance? *Perception & Psychophysics 68*: 800–814.

Schneider, K. A. (2011). Attention alters decision criteria but not appearance: A reanalysis of Anton-Erxleben, Abrams, and Carrasco (2010). *Journal of Vision 11*: 7. doi: 10.1167/11.13.7.

Schneider, K. A. and Komlos, M. (2008). Attention biases decisions but does not alter appearance. *Journal of Vision 8*: 3.1–10.

Schwartz, S., Vuilleumier, P., Hutton, C., Maravita, A., Dolan, R. J., and Driver, J. (2005). Attentional load and sensory competition in human vision: Modulation of fMRI responses by load at fixation during task-irrelevant stimulation in the peripheral visual field. *Cerebral Cortex 15*: 770–786.

Scolari, M., Byers, A., and Serences, J. T. (2012). Optimal deployment of attentional gain during fine discriminations. *Journal of Neuroscience 32*: 7723–7733.

Seidl, K. N., Peelen, M. V., and Kastner, S. (2012). Neural evidence for distractor suppression during visual search in real-world scenes. *Journal of Neuroscience 32*: 11812–11819.

Serences, J. T. and Boynton, G. M. (2007). Feature-based attentional modulations in the absence of direct visual stimulation. *Neuron 55*: 301–312.

Serences, J. T. and Saproo, S. (2012). Computational advances towards linking BOLD and behavior. *Neuropsychologia 50*: 435–446.

Serences, J. T., Saproo, S., Scolari, M., Ho, T., and Muftuler, L. T. (2009). Estimating the influence of attention on population codes in human visual cortex using voxel-based tuning functions. *NeuroImage 44*: 223–231.

Serences, J. T., Schwarzbach, J., Courtney, S. M., Golay, X., and Yantis, S. (2004). Control of object-based attention in human cortex. *Cerebral Cortex 14*: 1346–1357.

Serences, J. T., Yantis, S., Culberson, A., and Awh, E. (2004). Preparatory activity in visual cortex indexes distractor suppression during covert spatial orienting. *Journal of Neurophysiology 92*: 3538–3545.

Seung, H. S. and Sompolinsky, H. (1993). Simple models for reading neuronal population codes. *Proceedings of the National Academy of Sciences USA 90*: 10749–10753.

Shadlen, M. N., Britten, K. H., Newsome, W. T., and Movshon, J. A. (1996). A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *Journal of Neuroscience 16*: 1486–1510.

Shadlen, M. N. and Newsome, W. T. (1994). Noise, neural codes and cortical organization. *Current Opinion in Neurobiology 4*: 569–579.

Shannon, S. E. (1949). *A Mathematical Theory of Communication*. Urbana, Ill. University of Illinois Press.

Shaw, M. L. (1984). Division of attention among spatial locations: A fundamental difference between detection of letters and detection of luminance increments. In H. Bouma and D. G. Bouwhais (eds.), *Attention and Performance* (vol. 10, pp. 109–121). Hillsdale, N.J.: Lawrence Erlbaum Associates.

Shomstein, S., Kimchi, R., Hammer, M., and Behrmann, M. (2010). Perceptual grouping operates independently of attentional selection: Evidence from hemispatial neglect. *Attention, Perception, & Psychophysics 72*: 607–618.

Slotnick, S. D., Hopfinger, J. B., Klein, S. A., and Sutter, E. E. (2002). Darkness beyond the light: Attentional inhibition surrounding the classic spotlight. *NeuroReport 13*: 773–778.

Somers, D. C., Dale, A. M., Seiffert, A. E., and Tootell, R. B. (1999). Functional MRI reveals spatially specific attentional modulation in human primary visual cortex. *Proceedings of the National Academy of Sciences USA 96*: 1663–1668.

Stokes, M., Thompson, R., Nobre, A. C., and Duncan, J. (2009). Shape-specific preparatory activity mediates attention to targets in human visual cortex. *Proceedings of the National Academy of Sciences USA 106*: 19569–19574.

Sundberg, K. A., Mitchell, J. F., Gawne, T. J., and Reynolds, J. H. (2012). Attention influences single unit and local field potential response latencies in visual cortical area V4. *Journal of Neuroscience 32*: 16040–16050.

Sundberg, K. A., Mitchell, J. F., and Reynolds, J. H. (2009). Spatial attention modulates center-surround interactions in macaque visual area V4. *Neuron 61*: 952–963.

Sylvester, C. M., Shulman, G. L., Jack, A. I., and Corbetta, M. (2009). Anticipatory and stimulus-evoked blood oxygenation level-dependent modulations related to spatial attention reflect a common additive signal. *Journal of Neuroscience 29*: 10671–10682.

Teller, D. Y. (1984). Linking propositions. *Vision Research 24*: 1233–1246.

Tipper, S. P. and Driver, J. (1988). Negative priming between pictures and words in a selective attention task: Evidence for semantic processing of ignored stimuli. *Memory & Cognition 16*: 64–70.

Treisman, A. M. (1960). Contextual cues in selective listening. *Quarterly Journal of Experimental Psychology 12*: 242–248.

Treisman, A. M. and Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology 12*: 97–136.

Treue, S. (2003). Climbing the cortical ladder from sensation to perception. *Trends in Cognitive Sciences 7*: 469–471.

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

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

Van der Burg, E., Talsma, D., Olivers, C. N., Hickey, C., and Theeuwes, J. (2011). Early multisensory interactions affect the competition among multiple visual objects. *NeuroImage 55*: 1208–1218.

Vogel, E. K., Luck, S. J., and Shapiro, K. L. (1998). Electrophysiological evidence for a post-perceptual locus of suppression during the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance 24*: 1656–1674.

Vogel, E. K., Woodman, G. F., and Luck, S. J. (2005). Pushing around the locus of selection: Evidence for the flexible-selection hypothesis. *Journal of Cognitive Neuroscience 17*: 1907–1922.

Vuilleumier, P., Sagiv, N., Hazeltine, E., Poldrack, R. A., Swick, D., Rafal, R. D., and Gabrieli, J. D. (2001). Neural fate of seen and unseen faces in visuospatial neglect: A combined event-related functional MRI and event-related potential study. *Proceedings of the National Academy of Sciences USA 98*: 3495–3500.

Wijers, A. A., Lange, J. J., Mulder, G., and Mulder, L. J. (1997). An ERP study of visual spatial attention and letter target detection for isoluminant and nonisoluminant stimuli. *Psychophysiology 34*: 553–565.

Wojciulik, E., Kanwisher, N., and Driver, J. (1998). Covert visual attention modulates face-specific activity in the human fusiform gyrus: fMRI study. *Journal of Neurophysiology 79*: 1574–1578.

Yantis, S. and Johnston, J. C. (1990). On the locus of visual selection: Evidence from focused attention tasks. *Journal of Experimental Psychology: Human Perception and Performance 16*: 135–149.

Yigit-Elliott, S., Palmer, J., and Moore, C. M. (2011). Distinguishing blocking from attenuation in visual selective attention. *Psychological Science: A Journal of the American Psychological Society/APS 22*: 771–780.

Zenon, A. and Krauzlis, R. J. (2012). Attention deficits without cortical neuronal deficits. *Nature 489*: 434–437.

Zhang, W. and Luck, S. J. (2009). Feature-based attention modulates feedforward visual processing. *Nature Neuroscience 12*: 24–25.

Zhang, X., Zhaoping, L., Zhou, T., and Fang, F. (2012). Neural activities in V1 create a bottom-up saliency map. *Neuron 73*: 183–192.