CHAPTER 9

Using biologically plausible neural models to specify the functional and neural mechanisms of visual search

Glyn W. Humphreys*, Harriet A. Allen and Eirini Mavritsaki

Behavioural Brain Sciences, School of Psychology, University of Birmingham, Birmingham B15 2TT, UK

Abstract: We review research from our laboratory that attempts to pull apart the functional and neural mechanisms of visual search using converging, inter-disciplinary evidence from experimental studies with normal participants, neuropsychological studies with brain lesioned patients, functional brain imaging and computational modelling. The work suggests that search is determined by excitatory mechanisms that support the selection of target stimuli, and inhibitory mechanisms that suppress irrelevant distractors. These mechanisms operate through separable though overlapping neural circuits which can be functionally decomposed by imposing model-based analyses on brain imaging data. The chapter highlights the need for inter-disciplinary research for understanding complex cognitive processes at several levels.

Introduction

The visual world presents us with a complex and dynamically changing environment where it is important to be able to select efficiently stimuli that match our current behavioural goals. Attempts to measure the efficiency of visual selection have frequently used visual search to examine which factors facilitate and which impair selection processes. Across the past 40 years, numerous studies of search have been conducted which have demonstrated that, when targets differ from distractors in terms of their basic features (their colour, shape, size and so forth), then search is relatively efficient. In such cases the time taken to find the target increases by less than 10 ms/item as the number of distractors increases. In contrast, when targets and distractors share features, then search is much less efficient, with target detection times increasing often by 30 ms or more for each distractor present (see Wolfe, 1994). These contrasting patterns of search have often been characterised in terms of a two-stage account of visual selection (e.g. Neisser, 1967). According to this two-stage account, there is a first pre-attentive stage of visual processing which operates in parallel across the visual field and codes simple visual features. Targets that differ from distractors in their coding at this stage (having different features) can be detected in a spatially parallel manner. Targets that share features with distractors will activate overlapping representations at the pre-attentive stage and require further processing before they can be detected. This further processing is carried out at the second, attentive stage

---

*Corresponding author.

where there is serial scrutiny of each item — often this is conceived in terms of a serial window of attention being shifted from item to item. Due to this serial scrutiny, search rates increase to at least a minimal time required to make serial shifts of attention, and search time can be linearly related to the number of distractors present. Two-stage theories of this type remain highly influential. For example, Treisman's Feature Integration Theory (FIT; Treisman, 1998; Treisman and Gelade, 1980) maintains a distinction between feature-based pre-attentive stages and an attentional stage required to conjoin features. Wolfe's Guided Search Theory (GST) also proposes a first stage where simple features are coded and a second stage in which the items signalled as being most different from their neighbours are serially selected (Wolfe, 1994).

Other accounts, however, maintain that, rather than there being a strict dichotomy between pre-attentive and attentive stages of vision, there is a continuum of search efficiencies determined by different relations between targets and distractors. In their 'Attentional Engagement Theory' (AET), Duncan and Humphreys (1989, 1992) proposed that search efficiency was determined by the similarity between the target and the distractors (as above), and also by the similarity between distractors. High target-distractor similarity hinders search efficiency. High distractor similarity, on the other hand, can facilitate search because it enables distractors to be grouped and segmented from the target. Thus, even when targets and distractors share features search can be efficient if the distractors are homogeneous and can be grouped and rejected together (see also Humphreys et al., 1989). Duncan and Humphreys (1989, 1992) proposed that distractors were rejected together by a process of 'spreading suppression' when they grouped separate from targets. Quite similar ideas have subsequently been incorporated into traditional two-stage accounts in order to explain effects of distractor homogeneity. For example, FIT assumes that a process of distractor inhibition can be recruited which enables distractors to be rejected en masse through suppression of their common feature(s) (Treisman and Sato, 1990). In GST, distractor suppression operates through a process of lateral inhibition, so that distractors with common features will tend to suppress one another.

In addition to the process of rejecting distractors, accounts such as GST assume that search is guided by top-down activation of target features. This gives known targets a competitive edge in search tasks, enabling them rather than distractors to be selected. In terms of AET, targets are given a competitive advantage due to their having a template held in working memory which has a higher resting activation value than any template for distractors (see also Bundesen, 1990, for a similar idea expressed in terms of the target having a higher 'pertinence value'). There is clear behavioural evidence that having foreknowledge of the target makes a large difference to search, even determining whether stimuli 'pop out' or not. For example, large targets can be detected efficiently in the presence of small and medium-sized distractors, but medium targets are detected inefficiently amongst large and small distractors. Rather than this solely reflecting a difference in coding at the first pre-attentive stage, Hodsoll and Humphreys (2001) showed that efficient search depended on foreknowledge of the target. Having a template for the target enabled the feature differences to be used to guide search efficiently.

Modelling search

These ideas of search being guided by top-down activation from a target template, and also by distractor suppression, can be incorporated into more formal accounts of search including both mathematical (Bundesen, 1990) and computational models (e.g. Heinke and Humphreys, 2003; Humphreys and Müller, 1993; Mavritsaki et al., 2006). One value of such models is that, by demonstrating whether a proposed architecture can generate plausible search results, they provide an existence proof that the mechanisms of search could operate in the manner proposed. For example, since linear search functions can be generated by processes operating in a spatially parallel manner, the models demonstrate that serial processing operations are not *necessarily* required. The models also provide ways of

analysing how complex processes interact to generate 'whole-system' behaviour, something that is otherwise difficult to specify. Furthermore, if the models can incorporate processes that mimic real neuronal firing, then further physiological constraints can be added to the constraints of having to capture a body of psychological data, to given a multi-level account of human performance. This is the approach we have tried to follow when implementing the sSoTS (spiking Search over Time and Space) model (Mavritsaki et al., 2006), shown in Fig. 1.

sSoTS incorporates processes proposed by nearly all the major psychological accounts of search. Within the model visual input is coded by activating topographic maps representing simple visual features. Within each map the units interact through lateral inhibition, enhancing activation for a stimulus that differs from its local neighbours. This activation is transmitted to a 'master map' that sums activity for a given location within each feature map and then feeds back this activation to 'sharpen' the competition for selection at the feature map level. In addition, top-down activation is transmitted to the feature maps, both to increase the activation for target features and to decrease activity for known distractor features. This top-down activation can give the target a competitive edge enabling it to be selected ahead of the distractors, with target selection determined by setting a threshold for units within the master map of locations. Search efficiency in the model is determined by the overlap in features shared by targets and distractors and by whether the distractors have common or different features. Search operates in a spatially parallel manner across all of the items present, with efficiency decreasing linearly as targets and distractors share features. Activation profiles for the model are shown in Fig. 2 for two 'standard' search cases: A — where the target is defined by a difference in a single feature (SF) relative to distractors (e.g. target = blue H (*italic*) amongst blue As), and B — target = a conjunction of features, each of which is shared with a distractor (target = blue H and distractors blue As and green Hs). Activation in units corresponding to the location of the target rise more rapidly in the single feature relative to

the conjunction case. Figure 3a shows the RTs for units at the target location to reach a critical threshold point (signalling a target detection response) in the single feature and conjunction (CJ) conditions as a function of the number of distractors present. Search is more efficient (increasing less with the display size) in the single feature than the conjunction case. These data capture the difference in search efficiency between single feature and conjunction search (see Fig. 3b). For illustration, data are also shown for a preview condition which uses the same items as conjunction search but presents one set of distractors for 800 ms prior to the onset of the second set of distractors, plus the target when present (see below for further discussion).

Models such as sSoTS cannot only integrate different psychological proposals but they can also generate predictions about how processing may operate at a neural level, given that the model is based on the operation of biologically plausible processing units (simulate spiking neurons). For example, the feature maps may plausibly be located within areas of early visual cortex which respond to simple visual features. However, the master map is more likely to be located within posterior parietal cortex where there is evidence for neurons to be coded to the locations but not necessarily the features of visual stimuli (Courtney et al., 1994). If units within this location map are damaged, sSoTS predicts that search efficiency will deteriorate, particularly for targets that share features with distractors. This occurs because the model is less able to 'sharpen' any competition for selection, particularly for units on the affected ('contralesional') side. The net result is that targets on that side become difficult to detect, particularly when they share features with distractors. This is illustrated in the predictions for both reaction times to detect the target shown in Fig. 4a, and in the predicted error rates (target misses), shown in Fig. 4b.

The units within the model also operate using time parameters mimicking those of real neurons. For example, after firing, units build up a calcium parameter which reduces the future likelihood of firing for a period — units enter a refractory state. The emergent dynamics of activity lead to clear

138
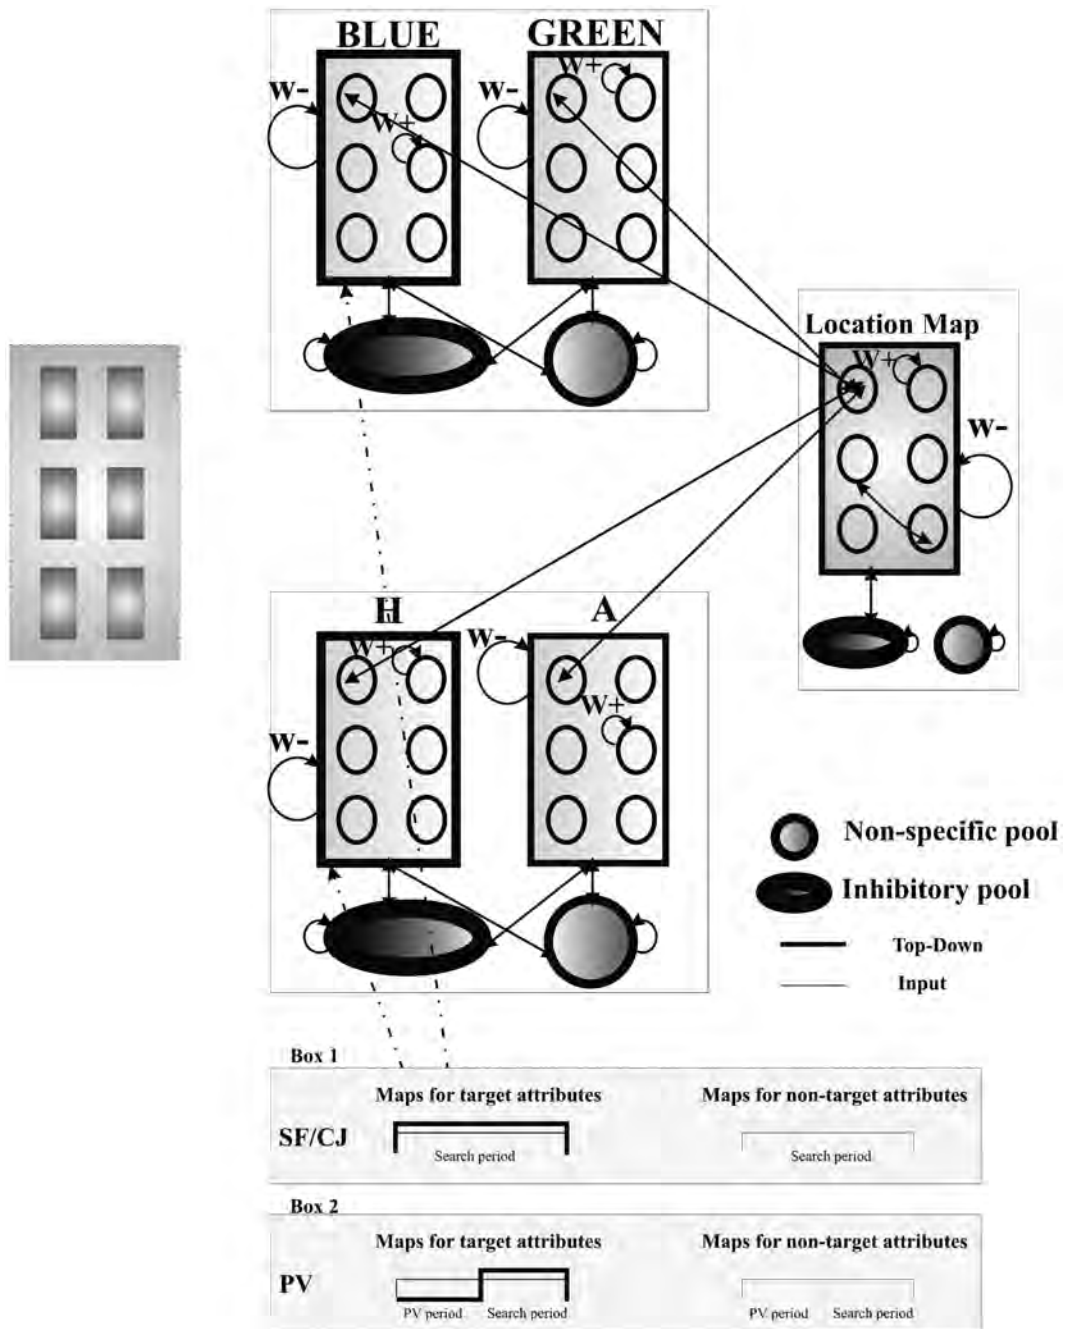
Fig. 1. The architecture of the sSoTS model. Input into the model is fed into the feature maps and from there into the location units. Units within maps, and at the same location across maps, operate in a mutually inhibitory way through the pools of inhibitory units. Activity in the location maps is fed back to the earlier maps, to bias competition for selection in favour of features that differ from their neighbours. In addition to such bottom-up biases, both excitatory and inhibitory activity can be set in a top-down manner, to facilitate target selection. Top-down excitation and inhibition helped to bias search to favour the target.
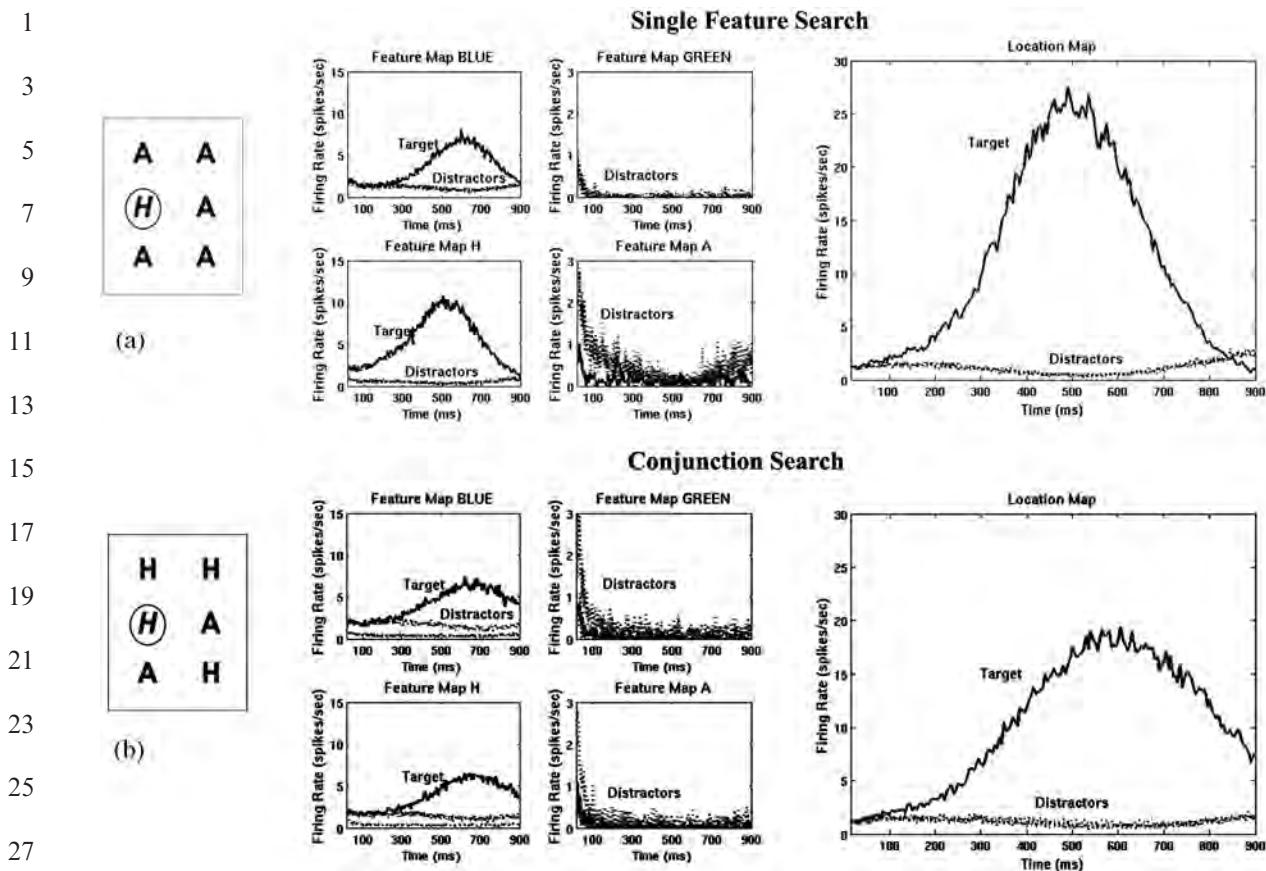
Fig. 2. (a) Activity in sSoTS plotted for four feature maps (BLUE, GREEN, H and A) and the Location map, when the target was a blue H (*italic*) and the distractors blue As. (b) Activity in the same maps for a conjunction target (blue H target, blue A and green H distractors). Activation in the Location map rises less rapidly and reaches a lower peak.

predictions being made about what might happen when the presentation of distractors is staggered over time. If there is presentation of one set of distractors prior to the other items, then activation for the initial distractors may be in a refractory state when the new items appear. Targets should be detected efficiently if they are presented at the time when distractor units are refractory, even if the distractors share their features with targets. This is illustrated in Fig. 5a.

Visual search over time

Predictions about the dynamics of visual search have been tested in studies using the 'preview'

search procedure. In this procedure the presentation of distractors is staggered over time, with one set of distractors appearing before the others and the target. This staggered presentation can facilitate target selection, as illustrated in Fig. 3b, where search is shown to be as efficient in the preview condition as in the single feature condition when only the new set of items is presented. sSoTS makes a matching prediction (Figs. 3a, b). Interestingly, and again like the model, there is a distinct time course to this effect — the first set of distractors needs to be presented up to 400 ms before the other items for search to benefit (Watson and Humphreys, 1997; Watson et al., 2003). This is a striking result because the different sets of stimuli can be temporarily segmented

AU : 4
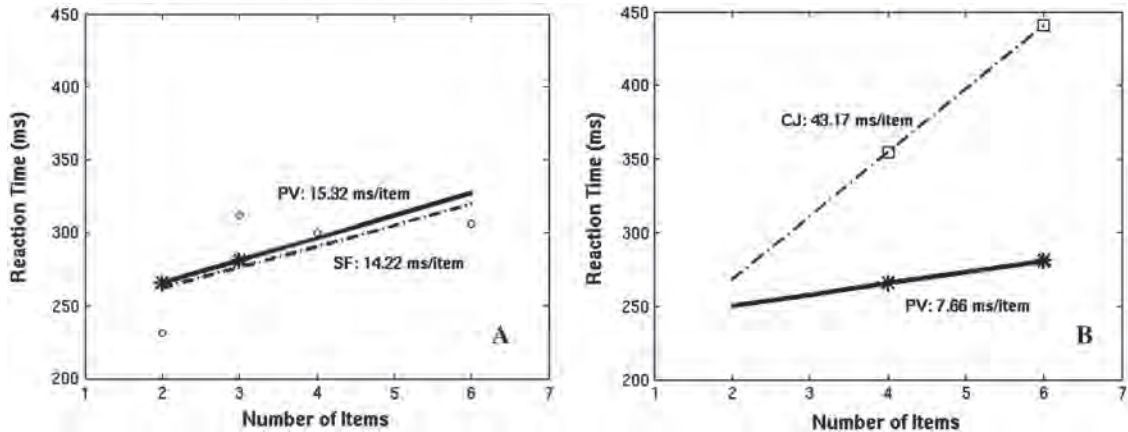
(a) Simulations

(b) Comparable human data (from Watson, Olivers & Humphreys, 2003)
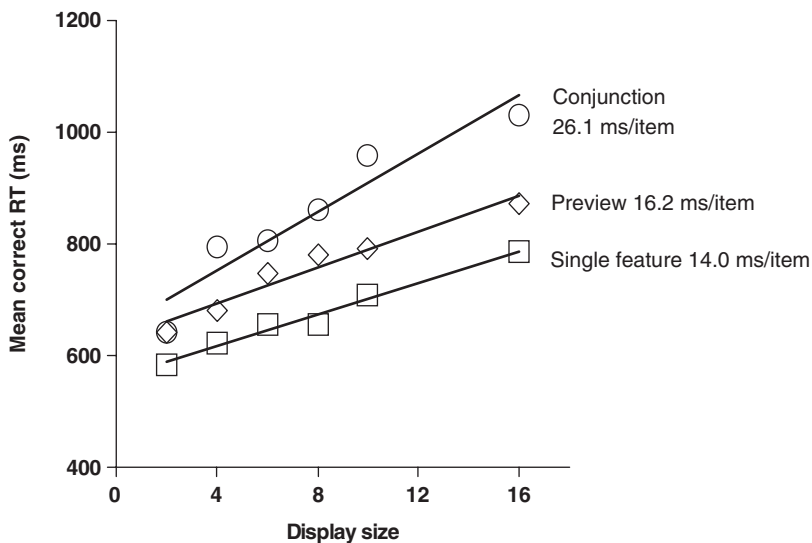
Fig. 3. (a) Simulations of single feature (SF), conjunction (CJ) and preview search (PV) is sSoTS. Note the steep slope on conjunction search even though the model operates in a spatially parallel manner. (b) Comparable human data (adapted from Watson et al., 2003). In human preview search, slopes for the preview condition very often match those in the single feature baseline (equivalent to when only the new search items are presented) and both are faster than the conjunction condition.

over much shorter time intervals than this — one can see that the old and new displays differ in time, but, with a short interval, it remains difficult not to be affected by the old items. This indicates that temporal segmentation alone is not sufficient to explain performance. However, the time course does match that expressed by sSoTS

(see Figs. 5a, b) — an emergent property of sSoTS's biologically plausible assumptions about the time course of the refractory state for neurons.

Other data from preview search also match other aspects of sSoTS. As we have noted, sSoTS incorporates the proposals made by psychological models that search is contingent on top-down
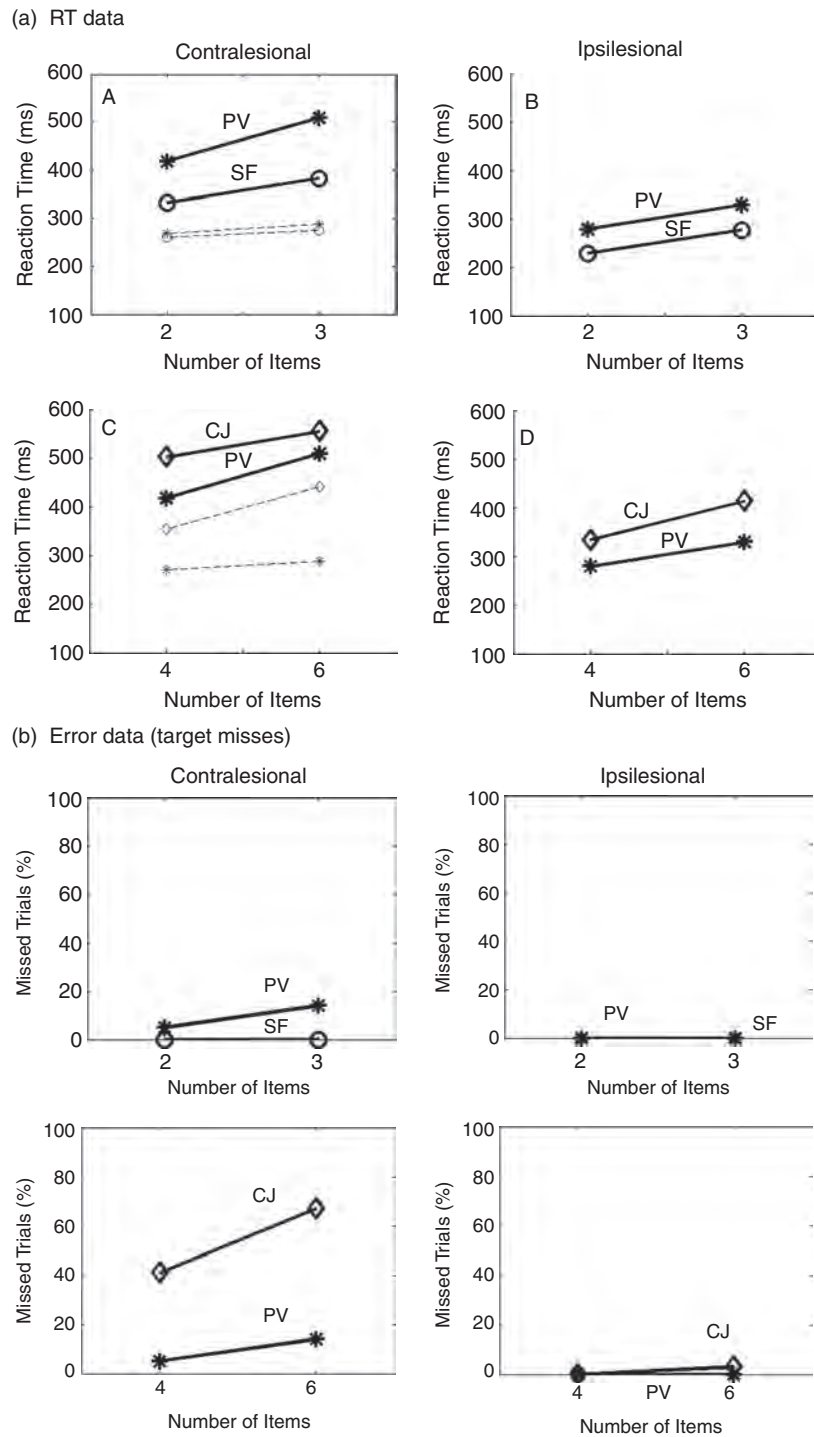
(a) RT data

Fig. 4. Simulations of the effects of lesioning units on one side of the location map in sSoTS. (a) RT data; (b) error data (target misses). The data are shown for targets falling on the contralesional or ipsilesional side of space (side affected by lesion; side unaffected by lesion). In the top figures, the data are shown for single feature (SF) and preview search (PV) according to the number of items in the final display of the preview condition (either two or three items). In the bottom figures the data are shown for the conjunction and preview conditions, plotted against the number of items in the final display (four or six, in both preview and conjunction search). In each case the dotted lines show the results when the model was unlesioned. The results indicate that lesioning disrupts search in the conjunction and preview conditions most, for contralesional targets.

142

(a) Predictions from sSoTS

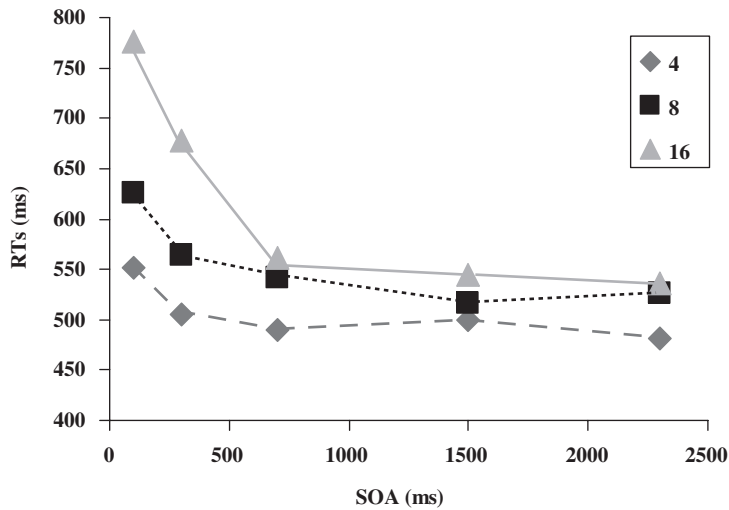Time Course VM: Experiment 1, Active Inhibition: 0.16
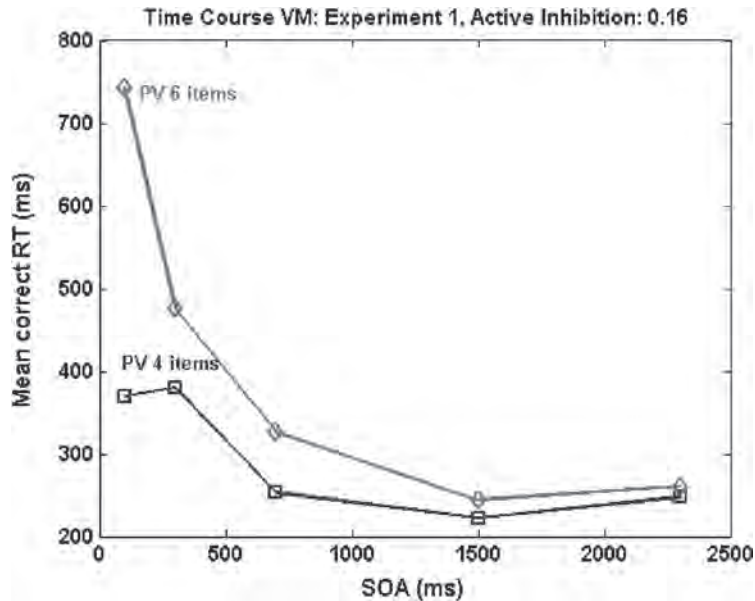
(b)

Fig. 5. (a) RTs generated by sSoTS as the duration of the preview is varied prior to the search display. (b) Data from human search as the preview duration is varied (adapted from Humphreys et al., 2004b). The model simulates the slow time course found in studies of preview search.

activation for targets and inhibitory suppression of distractors. Preview search provides good evidence for both processes operating in search. There are at least two pieces of evidence pointing to a role of inhibition. One comes from studies using probe-dot detection to measure where attention is allocated during search. In Humphreys et al's. (2004) study participants saw a set of distractors as a preview (e.g. green horizontal lines) followed by a search display (red vertical distractors and a red horizontal target, when present). On a majority of trials participants

AU : 5

carried out the search task. On a minority of trials they were cued to stop search and to try and detect a small probe that could appear either within an old distractor, a new distractor or the background. When the preview was presented for 800 ms before the new items, search for the new target was efficient. However, probes that fell on old items were difficult to detect, with detection levels in this case being lower than those found for probes presented on the background (see also Agter and Donk, 2005; Olivers and Humphreys, 2002; Watson and Humphreys, 1998). This is consistent with the spatial locations of the old items being inhibited.

A second piece of evidence indicating that there is inhibition of distractors comes from work on 'carry-over effects' in preview search. Braithwaite and Humphreys (2003) and Olivers and Humphreys (2003) presented a preview display of distractors in one colour followed by targets that either did or did not carry the colour of the to-be-ignored old distractors. Targets carrying the colour of the old distractor were difficult to detect — strikingly, this occurred even when the target had a singleton colour relative to the other new items being presented, and even when the old items were removed at the onset of the new displays. Normally such a colour single-ton should pop out in search. The problem in detecting such a singleton target provides strong evidence against the view that preview search is simply based on automatic detection of the new items or on the temporal segmentation of the old and new displays (cf. Donk and Theeuwes, 2001; Jiang et al., 2002) — if that were the case, then the singleton should have popped out. This negative colour carry-over effect is consistent with the inhibition of the features as well as the locations of the old items (cf. Humphreys et al., 2004). The result also fits with the idea of spreading suppression, as put forward by Duncan and Humphreys (1989); in this case, there is a spread of suppression from inhibited old distractors to new items carrying the inhibited properties — the result is that reaction times are slowed targets with these properties. This inhibition is maintained for at least some period even after the old distractors have been removed.

In addition to presenting evidence for the inhibition of old distractors, Braithwaite and Humphreys (2003) also reported data indicating effects of a positive expectancy for targets. In particular, Braithwaite and Humphreys showed that the negative colour carry-over effect could be reduced if participants had advanced knowledge of the target's colour. These authors propose that participants can independently set a top-down positive expectancy for a target along with adopting a negative bias against the properties of irrelevant distractors. The data indicating both positive and negative top-down effects in search match the top-down excitatory and inhibitory components operating in sSoTS.

The neural basis of inhibitory and excitatory biases

When people engage in visual search a range of brain areas are very often activated, most notably there is a conjunction of activity in posterior parietal and frontal cortices which increases as search becomes more difficult (see Corbetta and Shulman, 2002). However, as we have noted, search involves multiple processes (positive activation for targets, inhibitory suppression of distractors, the maintenance of target templates and so forth), so it is useful to explore paradigms such as preview search which can enable different processes to be isolated. There have now been several studies of preview search using functional brain imaging, and it has been consistently found that, relative to search when all the items appear together, preview search is associated with increased activation of several regions of posterior parietal cortex (the superior parietal lobe [SPL] and the precuneus; Allen et al., 2008; Humphreys et al., 2004; Olivers et al., 2005; Pollmann et al., 2003). This is interesting because preview search can be more efficient than baseline search conditions when all the items appear simultaneously, so the increased activation does not simply reflect the general difficulty of search. Both Allen et al. (2008) and Pollmann et al. (2003) also included some 'dummy preview' trials where only the preview appeared although participants

expected a search display to follow the preview (and so participants should engage in the same processing of the preview as on search trials). In both studies there was increased activation of SPL and precuneus when it was a dummy preview trial compared to trials which used equivalent visual displays but where the previews were unlikely to be ignored. This indicates that the SPL/precuneus activation is not tied to the search operation but it is consistent with these brain regions being linked to the inhibitory processing of distractors. The activation of the SPL/precuneus may reflect the operations of inhibitory neurons or some initial attention being paid to the old distractors in order to then inhibit them (see Humphreys et al., 2004, for evidence consistent with this from probe-dot procedures).

The data on functional brain imaging are supported by neuropsychological studies on selective disorders of search in patients with brain lesions. It is well established that patients with damage to posterior parietal cortex (PPC) are impaired at serial search tasks (Eglin et al., 1989; Riddoch and Humphreys, 1987). Olivers and Humphreys (2004) found that PPC patients were also impaired at preview search, even though normal participants perform preview search efficiently. This again points to effects on particular processes rather there being an exaggerated influence of search difficulty in the patients. Humphreys et al. (2006) found that PPC patients impaired at preview search were nevertheless able to prioritize their attention to new onset stimuli. If prioritized attention to new onsets was sufficient to generate efficient preview search (Donk and Theeuwes, 2001), then the patients should have shown efficient preview search. In contrast to this the impairments indicate that additional processes (such as inhibition of the old distractors) determine the efficiency of preview search, and these additional processes may be disrupted by PPC damage. Further evidence consistent with this was noted by Olivers and Humphreys (2004). They found that the PPC patients were particularly impaired when the new and old stimuli were spatially overlapping compared with when the items did not overlap and could be spatially segmented. These data suggest that the PPC was critical for the segmentation and rejection of the old distractors, and was recruited particularly when spatial segmentation was difficult.

As we have noted (Fig. 4), sSoTS also predicts that damage to the location units within the model, putatively representing PPC, generates problems in search. Within the model this not only affects the detection of conjunction targets, but also the detection of targets in preview search. Humphreys et al. (2009) simulated effects of PPC damage by removing units from one side of the location map. This disrupted both conjunction and preview search. Similarly to the patients, the problems in preview search were most pronounced when the old and new items overlapped spatially, when temporal segmentation would normally augment spatial segmentation (Fig. 6). When lesioned, sSoTS has a reduced ability to implement temporal separation of new targets from old distractors on one side of space; this leads to problems in separating distractors within the same area of field. It is of interest that poor performance of the model was observed even when both the old and new items fell within the undamaged (ipsilesional) field. This is because poor inhibition of old ipsilesional distractors means that they stay as competitors for new items appearing in the same locations. The result is that target detection becomes inefficient.

These simulations of the effects of PPC damage provide important converging evidence linking specific brain regions to functional processes within the model.

Modelling brain imaging data

Another way to link the operation of an abstract computational model to brain function is to simulate data from functional imaging. Activity within a model such as sSoTS can be convolved with an assumed haemodynamic function to predict the BOLD response (Glover, 1999). This then enables us to take functional processes in the model, such as the operation of excitatory and inhibitory activity during search, and to assess within which brain areas the activity correlates with the different functional processes. We

(a) Stimulus presentation procedure
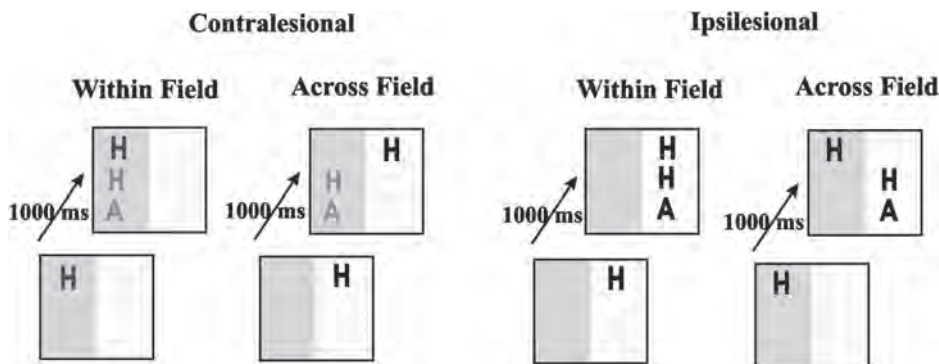
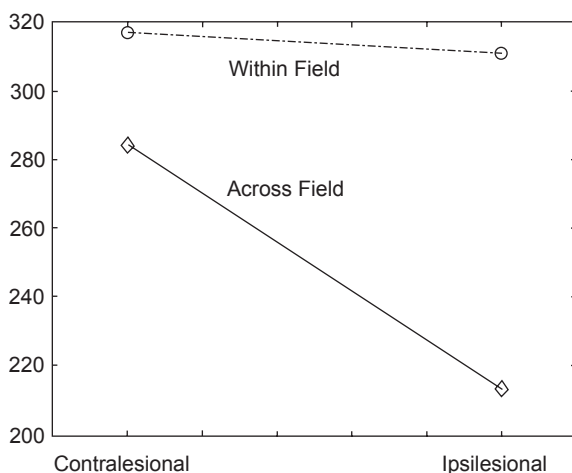
(b) Results (search RTs) from sSoTS after lesioning

Fig. 6. (a) The stimulus presentation procedure used to examine the effects of spatial and temporal separation in sSoTS. The initial display is the preview and this is followed by the search display. In the across field condition, the preview appears in the opposite field to the new items in the search display. In the within field condition, the new search items appear within the same field as the preview. (b) Mean RTs generated by sSoTS after unilateral lesioning of units on one side of the Location map. Targets were presented either on the contralesional (damaged units) or ipsilesional side of space (undamaged units). The data show that within field discriminations are more difficult than across field discriminations, with the detection of contralesional targets being generally worse. Comparable data from patients with PPC lesions were reported by Olivers and Humphreys (2004).

(Mavritsaki et al., submitted) have done this by summing activity across different maps in sSoTS according to whether the activity is reflecting top-down excitation of targets or bottom-up inhibition of distractors. The emergent results are shown in Fig. 7. The figure reveals both overlapping and distinct regions within PPC, and also visual processing areas in occipital cortex, where activity is separately correlated with the time course of

excitatory and inhibitory activity within the model. This then provides a functional decomposition of the network of areas that is activated during search. As the functional processes of target excitation and distractor inhibition will be involved both in preview search and in search for conjunction stimuli, it follows that damage to the areas supporting the excitation and inhibition processes will disrupt both types of search task.
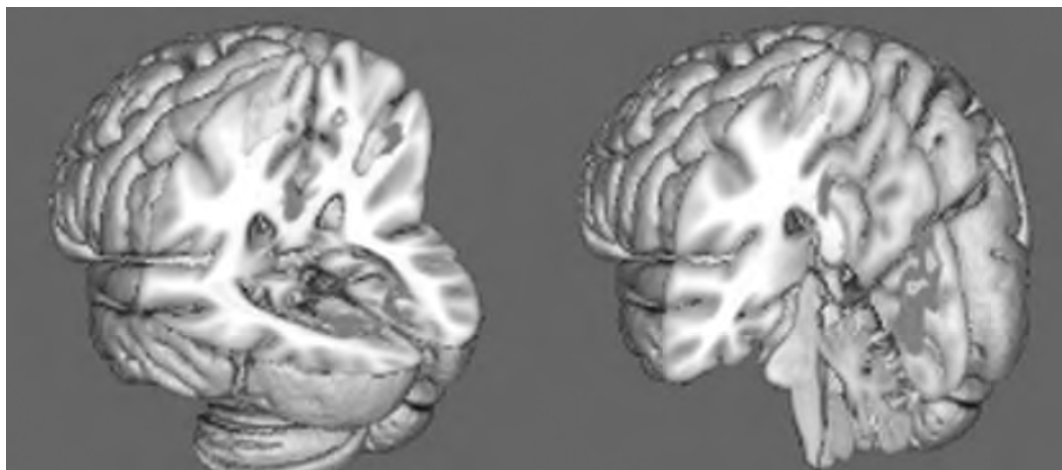
Fig. 7. Images showing brain regions where neural activity in preview search correlates with inhibitory and excitatory activity in sSoTS (when convolved with an assumed haemodynamic function). This model-based analysis pulls apart neurons within overlapping brain regions that perform functionally distinct roles in search. Areas in blue correlate with inhibitory activity in the model; areas in yellow correlate with excitatory activity.

This prediction is supported by the neuropsychological data (Olivers and Humphreys, 2004; Riddoch and Humphreys, 1987).

The importance of multi-level analyses

We have reviewed evidence using behavioural manipulations (such as preview search), neuropsychological analyses (e.g. the effects of PPC lesions), functional brain imaging (e.g. fMRI) and computational modelling, to analyse the processes involved in visual search. Each piece of evidence has its own limitation. For example, behavioural studies reveal 'whole-system' behaviour, but this can make it difficult to analyse the operation of sub-component processes. Studies using functional brain imaging reveal brain areas that correlate with different processes but do not prove that these processes are necessary for a given task. Neuropsychological studies do demonstrate the necessary role of brain regions (since damage to those regions is shown to disrupt performance), but the lesion may affect more than one process which in turn makes it difficult to exactly relate an impaired function to the lesioned area. Given these limitations, it is important to use

evidence coming from each approach in order to develop an over-arching framework not subject to one limitation. In the present case, this framework can also be captured at a formal level in terms of the sSoTS model. Models such as sSoTS can help to integrate research in at least two ways. First, it can simulate effects at multiple levels (brain imaging, effects of neuropsychological deficits and effects due to emergent, whole-system behaviour), enabling us to link the different approaches and different types of data together. The weakness inherent in one approach, then, can be compensated for by the strengths of the other. For example, fMRI in normal participants may be able to localise, across the whole brain, processes involved in a given task. The necessary role of these areas would then be addressed through neuropsychological evidence. This relationship should be captured by simulating the effects of lesioning matching regions within the model. Second, the model shows how ideas expressed in different psychological models can be formally linked, enabling us to see how different models relate to one another. For example, distractor similarity influences the amounts of lateral inhibition operating in the model (cf. Duncan and Humphreys, 1989; Wolfe, 1994), while Treisman

and Sato's argument for feature-based inhibition is implemented through the top-down inhibition process in the model. Which factors are important, under which conditions, can then be explored. We believe this provides a working framework through which to assess the various factors determining search efficiency. One particularly important point, given the stab at biological plausibility in the model, is that variations in physiological parameters (e.g. the time course over which neurons enter into a refractory state) can generate psychological predictions (e.g. on the time course of visual search); these predictions can be tested and fed back to further inform model development in a (virtual) cycle of modelling and testing. We propose that the formal development of models such as sSoTS will play an important part in the integration of psychological theory and physiological data.

Uncited reference

Humphreys et al. (2004a).

Acknolwedgement

References

Agter, , & Donk, M. (2005). Prioritized selection in visual search through onset capture and color inhibition: Evidence from a probe-dot detection-task. *Journal of Experimental Psychology: Human Perception and Performance*, *31*, 722–730.

Allen, H. A., Humphreys, G. W., & Matthews, P. M. (2008). A neural marker for content specific active ignoring. *Journal of Experimental Psychology: Human Perception and Performance*, *34*, 286–297.

Braithwaite, J. J., & Humphreys, G. W. (2003). Inhibition and anticipation in visual search: Evidence from effects of color foreknowledge on preview search. *Perception & Psychophysics*, *65*, 213–237.

Bundesen, C. (1990). A theory of visual-attention. *Psychological Review*, *97*(4), 523–547.

Corbetta, M., & Shulman, G. L. (2002). Control of goal directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, *3*, 201–215.

Courtney, S. M., Ungerleider, L. G., Keil, K., & Haxby, J. V. (1994). Object and spatial visual working memory activate separate neural systems in human cortex. *Cerebral Cortex*, *6*, 39–49.

Donk, M., & Theeuwes, J. (2001). Visual marking beside the mark: Prioritizing selection by abrupt onsets. *Perception & Psychophysics*, *93*, 891–900.

Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review*, *96*, 433–458.

Duncan, J., & Humphreys, G. W. (1992). Beyond the search surface: Visual search and attentional engagement. *Journal of Experimental Psychology: Human Perception and Performance*, *18*, 578–588.

Eglin, M., Robertson, L. C., & Rafal, R. D. (1989). Visual search performance in the neglect syndrome. *Journal of Cognitive Neuroscience*, *1*, 372–385.

Glover, G. (1999). Deconvolution of impulse response in event-related fMRI. *NeuroImage*, *9*, 416–429.

Heinke, D., & Humphreys, G. W. (2003). Attention, spatial representation and visual neglect: Simulating emergent attentional processes in the Selective Attention for Identification Model (SAIM). *Psychological Review*, *110*, 29–87.

Hodsoll, J., & Humphreys, G. W. (2001). Driving attention with the top down: The relative contribution of target templates to the linear separability effect in the size dimension. *Perception & Psychophysics*, *63*, 918–926.

Humphreys, G. W., Jung-Stalmann, B., & Olivers, C. N. L. (2004a). An analysis of the time course of attention in preview search. *Perception & Psychophysics*, *66*, 713–730.

Humphreys, G. W., Kyllingsbæk, S., Watson, D. G., Olivers, C. N. L., Law, I., & Paulson, O. (2004b). Parieto-occipital areas involved in efficient filtering in search: A time course analysis of visual marking using behavioural and functional imaging procedures. *Quarterly Journal of Experimental Psychology*, *57A*, 610–635.

Humphreys, G. W., Mavritsaki, E., Allen, H. A., Heinke, D., & Deco, G. (2009). Modelling visual search in biologically plausible neural networks: Whole-system behaviour, neuropsychological breakdown and BOLD signal activation. In D. Heinke & E. Mavritsaki (Eds.), *Computational modelling in behavioural neuroscience*. London: Psychology Press.

Humphreys, G. W., & Müller, H. M. (1993). SEarch via Recursive Rejection (SERR): A connectionist model of visual search. *Cognitive Psychology*, *25*, 43–110.

Humphreys, G. W., Olivers, C. N. L., & Yoon, E. Y. (2006). An onset advantage without a preview benefit: Neuropsychological evidence separating onset and preview effects in search. *Journal of Cognitive Neuroscience*, *18*, 110–120.

Humphreys, G. W., Quinlan, P. T., & Riddoch, M. J. (1989). Grouping effects in visual search: Effects with single- and combined-feature targets. *Journal of Experimental Psychology: General*, *118*, 258–279.

Jiang, Y., Marks, L. E., & Chun, M. M. (2002). Visual marking: Selective attention to asynchronous temporal groups. *Journal of Experimental Psychology: Human Perception and Performance*, *28*, 717–730.

148

Mavritsaki, E., Allen, H. A. & Humphreys, G. W. (submitted). Modelling the neural substrate of preview search. *Neuron*.

AU : 6

Mavritsaki, E., Heinke, D. G., Humphreys, G. W., & Deco, G. (2006). A computational model of visual marking using an interconnected network of spiking neurons: The spiking Search over Time & Space model (sSoTS). *Journal of Physiology (Paris)*, *100*, 110–124.

Neisser, U. (1967). *Cognitive psychology*. New York: Appleton-Century-Crofts.

Olivers, C., & Humphreys, G. W. (2002). When visual marking meets the attentional blink: More evidence for top-down limited capacity inhibition. *Journal of Experimental Psychology: Human Perception and Performance*, *28*, 22–42.

Olivers, C. N., & Humphreys, G. W. (2004). Spatiotemporal segregation in visual search: Evidence from parietal lesions. *Journal of Experimental Psychology: Human Perception and Performance*, *30*, 667–688.

Olivers, C. N. L., & Humphreys, G. W. (2003). Visual marking and singleton capture: Fractionating the unitary nature of visual selection. *Cognitive Psychology*, *47*, 1–42.

Olivers, C. N. L., Smith, S., Matthews, P., & Humphreys, G. W. (2005). Prioritizing new over old: An fMRI study of the preview search task. *Human Brain Mapping*, *24*, 69–78.

Pollmann, S., Weidner, R., Humphreys, G. W., Olivers, C. N. L., Müller, K., Lohmann, G., et al. (2003). Separating segmentation and target detection in posterior parietal cortex — An event-related fMRI study of visual marking. *NeuroImage*, *18*, 310–323.

Riddoch, M. J., & Humphreys, G. W. (1987). Perceptual and action systems in unilateral neglect. In M. Jeannerod (Ed.), *Neurophysiological and neuropsychological aspects of spatial neglect*. Amsterdam: Elsevier Science.

Treisman, A. (1998). Feature binding, attention and object perception. *Philosophical Transactions of the Royal Society*, *353*, 1295–1306.

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

Treisman, A., & Sato, S. (1990). Conjunction search revisited. *Journal of Experimental Psychology: Human Perception and Performance*, *16*, 459–478.

Watson, D. G., & Humphreys, G. W. (1997). Visual marking: Prioritising selection for new objects by top-down attentional inhibition. *Psychological Review*, *104*, 90–122.

Watson, D. G., & Humphreys, G. W. (1998). Visual marking of moving objects: A role for top-down feature based inhibition in selection. *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 946–962.

Wolfe, J. M. (1994). Guided Search 2.0: A revised model of visual search. *Psychonomic Bulletin and Review*, *1*, 202–238.