VU Research Portal

A Boost and Bounce theory of temporal attention

Olivers, C.N.L.; Meeter, M.

Link to publication in VU Research Portal

A Boost and Bounce Theory of Temporal Attention

Christian N. L. Olivers and Martijn Meeter
Vrije Universiteit Amsterdam

What is the time course of visual attention? Attentional blink studies have found that the 2nd of 2 targets is often missed when presented within about 500 ms from the 1st target, resulting in theories about relatively long-lasting capacity limitations or bottlenecks. Earlier studies, however, reported quite the opposite finding: Attention is transiently enhanced, rather than reduced, for several hundreds of milliseconds after a relevant event. The authors present a general theory, as well as a working computational model, that integrate these findings. There is no central role for capacity limitations or bottlenecks. Central is a rapidly responding gating system (or attentional filter) that seeks to enhance relevant and suppress irrelevant information. When items sufficiently match the target description, they elicit transient excitatory feedback activity (a "boost" function), meant to provide access to working memory. However, in the attentional blink task, the distractor after the target is accidentally boosted, resulting in subsequent strong inhibitory feedback response (a "bounce"), which, in effect, closes the gate to working memory. The theory explains many findings that are problematic for limited-capacity accounts, including a new experiment showing that the attentional blink can be postponed.

Keywords: attention, time course, attentional blink, awareness

One of the brain's crucial functions is to prioritize relevant over irrelevant information. It does this by a set of mechanisms we collectively call selective attention. By its very definition, selective attention is selective: It is thought that only one or, at most, a few objects can be processed at a time. The important question then is how much time is spent on selecting and processing one object or set of objects before attention is available again for the next selection. In other words, what are the dynamics of attention? What is its time course? Or, as others have phrased it, what is the "attentional dwell time" (Duncan, Ward, & Shapiro, 1994)?

The Answer From the Attentional Blink: Attention Is Slow and Deals With One Object at a Time

If one is interested in the pure time course of attention ("how long does it take before one can attend to the next piece of relevant information") one needs a paradigm that controls for other factors that may take time, such as having to switch from one location to the other, or from one complete task to the other. Probably the most popular paradigm in this respect is the rapid serial visual presentation task (RSVP; Lawrence, 1971). For the past 15 years, the two-targets version of the RSVP task has dominated the literature. Figure 1a depicts a typical example, in which participants are asked to report two letters from a stream of digits all presented at a single location (Chun & Potter, 1995). Participants have little difficulty reporting the first target (T1). However, as shown in Figure 1b, report of the second target (T2) suffers considerably when it is presented within about 500 ms from T1. It is as if attention blinks for half a second while it is busy processing T1; hence, the phenomenon has been termed the *attentional blink* (Raymond, Shapiro, & Arnell, 1992).

To our knowledge, all currently active theories of the attentional blink attribute it to a limited-capacity processing stage, or bottleneck, with a relatively late locus in the information-processing stream. The idea is that T1 temporarily uses up vital mental resources that are then not available to T2. Figure 1c illustrates this resource depletion. The prototypical account is Chun and Potter's (1995) *two-stage theory*, but similar proposals were made earlier by Broadbent and Broadbent (1987) and a little later by Jolicoeur and colleagues (Jolicoeur, 1998, 1999; Jolicoeur & Dell'Acqua, 1998; Jolicoeur, Tombu, Oriet, & Stevanovski, 2002). According to two-stage theories, all items in the stream can be processed up to semantic levels (often referred to as conceptual short-term memory), but representations are vulnerable. For conscious report of a target, a second processing stage is required which consolidates the item in short-term memory proper. This consolidation process is limited to about one item at a time, and, it is important to note, takes time. Thus, while T1 is being consolidated, T2 must wait. Waiting takes longer if T1 processing takes longer, for instance, when T1 is masked by a subsequent distractor. However,

A. Typical attentional blink task

B. Typical data (Chun & Potter, 1995)
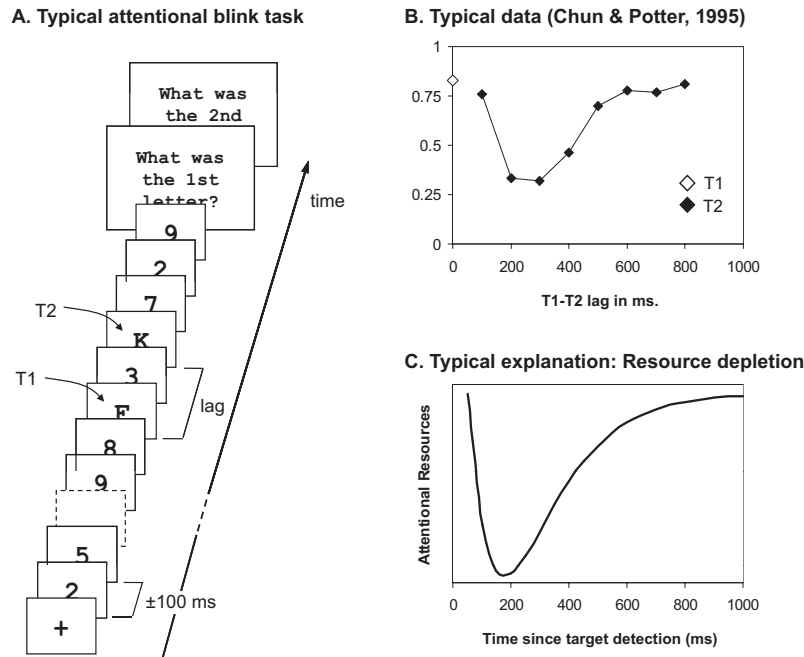
C. Typical explanation: Resource depletion

Figure 1. (A) The prototypical attentional blink task, in which the observer is asked to report the two letter targets (T1 and T2) embedded in a stream of digit distractors. The lag between the two targets is varied. (B) Typical performance for T1 and T2 as a function of lag (adapted from "A Two-Stage Model for Multiple Detection in Rapid Serial Visual Presentation," M. M. Chun & M. C. Potter, 1995, *Journal of Experimental Psychology: Human Perception and Performance, 21,* Figure 2, p. 112. Copyright 1995 by the American Psychological Association. (C) The classic explanation of the attentional blink: T1 processing leads to a depletion of resources that then cannot be used for T2.

there is little opportunity for waiting, because a distractor quickly follows T2, overwrites its vulnerable first-stage representation, and causes an attentional blink as a result. Thus, distractors play two important roles in two-stage theories: (a) They render T1 processing more difficult, resulting in longer waiting times and, thus, a more profound blink; and (b) they interfere with the T2 representation, such that T2 becomes sensitive to the temporary resource depletion (Brehaut, Enns, & Di Lollo, 1999; Chun & Potter, 1995; Giesbrecht & Di Lollo, 1998; Grandison, Ghirardelli, & Egeth, 1997; Seiffert & Di Lollo, 1997). Indeed, it has been shown that a skeletal version of the paradigm, with just the two targets and their immediately following distractors, is sufficient to generate an attentional blink (Ward, Duncan, & Shapiro, 1997).

A similar account has been proposed by Shapiro and Raymond, with their colleagues (Isaak, Shapiro, & Martin, 1999; Raymond, Shapiro, & Arnell, 1995; Shapiro & Raymond, 1994; Shapiro, Raymond, & Arnell, 1994). According to their *interference theory*, T1, T2, and the respective subsequent distractors are uploaded into visual short-term memory (VSTM). Within VSTM, these items then compete for conscious report. However, because VSTM resources are limited, the competition is biased in favor of those items that entered first, with further biases toward those items that closely resemble the targets. Thus, T1 and the distractor immediately following it often win the competition for retrieval, rather than T1 and T2. Within interference theory, the attentional blink period corresponds to the time needed for T1 to win the competition and be transferred to a report stage, so that

VSTM can be cleared for T2. As in two-stage theory then, items are initially represented in a feeble memory system (conceptual short-term memory or VSTM) and, at that moment, compete for limited resources. These resources are primarily assigned to T1 because it appears first. While T1 undergoes further processing, T2 suffers from either masking or competition for selection from surrounding distractors and is lost (see Shapiro, Arnell, & Raymond, 1997, for a unified model).

The *temporary loss of control theory* (TLC; Di Lollo, Kawahara, Ghorashi, & Enns, 2005; Kawahara, Kumada, & Di Lollo, 2006) has been presented as an alternative to limited-capacity accounts. It attributes the attentional blink to a temporary loss of control over the stimulus input. It assumes that, leading up to T1, the system is configured as an input filter: It adopts an attentional set for targets and against distractors. Targets pass the filter into higher processing stages necessary for consolidation and response planning, whereas distractors are rejected. An important assumption is that, because targets and distractors are usually arbitrarily defined, the input filter must be actively maintained by a central executive process. However, when T1 passes the filter, its consolidation also requires the dedication of the central executive. This means that the central executive can no longer maintain the attentional set, and the input filter becomes vulnerable to stimulus-driven disruption from distractors. As a consequence, targets are no longer allowed to enter, and an attentional blink is observed. As the central executive gradually becomes available again, the cor-

rect input filter is reinstated, and performance returns to normal. It is questionable, however, whether TLC indeed manages to avoid the limited-capacity resource-depletion argument. Notably, it assumes that T1 occupies a central executive for some time, during the course of which the system is not ready for T2. It appears then that limited-capacity resources have entered through the back door.

Several computational models have recently been developed that are based on the above ideas (Bowman & Wyble, 2007; Battye, 2006; Chartier, Cousineau, & Charbonneau, 2004; Dehaene, Sergent, & Changeux, 2003; Fragopanagos, Kockelkoren, & Taylor, 2005; Shih, 2008). As we discuss further in the General Discussion, they all assume either direct competition between T1 and T2 and/or indirect competition via the drainage of some limited-capacity resource. In summary then, the consensus is that the attentional blink reflects a relatively long-lasting limitation on attention after encountering a relevant visual event.

The Answer From Cueing Paradigms: Attention Is Fast and Can Deal With Multiple Relevant Items Presented in Rapid Sequence

Other studies have suggested an attentional time course that is rather opposite to that suggested by the attentional blink. Some of these are illustrated in Figure 2a. For example, Nakayama and Mackeben (1989) asked observers to detect and identify a target in a cluttered display, followed by a mask. Preceding the target display, a cue indicated the target position with 100% validity. Note that because cue, target, and mask always appeared in the same position, this procedure is not at all unlike RSVP. The results, however, were rather different: Target-identification performance increased with increasing stimulus onset asynchrony (SOA) between cue and target display, up to about 100–200 ms. Beyond this peak in performance, increasing SOAs resulted in gradually decreasing cueing benefits over the time course of several hundreds of milliseconds. Figure 2b (top left panel) shows Nakayama and Mackeben's findings, whereas Figure 2c illustrates the hypothesized steeply rising and slowly decaying underlying attentional time course function. Nakayama and Mackeben referred to this temporary enhancement in performance as *transient attention* and argued that it is largely automatic, operating at a relatively early level in visual processing. They distinguished it from a slower, more sustained attentional component. Similar transient attentional enhancement functions have been found or hypothesized by others using various paradigms, including cued visual search (Kristjánsson, Mackeben, & Nakayama, 2001), classic spatial cueing (Müller & Rabbitt, 1989), spatial distortion (Suzuki & Cavanagh, 1997), image-classification techniques (Shimozaki, Chen, Abbey, & Eckstein, 2007), saliency-guided search (Nothdurft, 2002), illusory line motion (Hikosaka, Miyauchi, & Shimojo, 1993), temporal order judgment (Scharlau, Ansorge, & Horstmann, 2006), flash-lag effects (Bachmann & Oja, 2003), eye-movement measurements (Mackeben & Nakayama, 1993), and RSVP (Botella, Barriopedro, & Suero, 2001; Chua, Goh, & Hon, 2001).

A rapid transient attention function was actually proposed earlier by Weichselgartner and Sperling (1987), using a combination of cueing and RSVP. In some of their experiments, observers monitored a single stream of digits, waiting for a particular cue (which was typically an outline square or highlighted digit). The task was to report the cued digit as well as the digits following it. They found that the cued digit and the one immediately following

it had a high likelihood of report, as had digits later in the stream (from 300 ms onward). In between, at about 200 ms, was a brief dip. Figure 2b (top right panel) shows this pattern of results. Like Nakayama and Mackeben (1989) did later, Weichselgartner and Sperling argued for an initial, rapid, and automatic attentional component that was then followed by a slower, more controlled component. Weichselgartner and Sperling referred to these two components as two attentional "glimpses." The brief dip at 200 ms represented the transition between these two glimpses.

It deserves mentioning that Raymond et al. (1992) argued that this dip might actually reflect an attentional blink. We agree, but, as we explain later, if it is an attentional blink, it is not caused by capacity limitations (nor was this claimed as such by Raymond et al., 1992). The typical limited-capacity explanation would be that the cued digit drains attentional resources, at the expense of the digit presented at 200 ms. However, there are several aspects of Weichselgartner and Sperling's (1987) data that go against this possibility. For one, the dip disappeared when a faint outline square was used to indicate the first to-be-reported digit, or when an auditory instead of visual cue was used. This suggests there may have been some low-level visual interference involved. More important is the fact that, within the second glimpse, observers were able to report more than the two targets so typical for the attentional blink task. For example, observers might report one digit in their "first glimpse" and then three subsequent digits in their "second glimpse" without much trouble. Limited-capacity theories would predict a second attentional blink in the second glimpse (see, e.g., Chun & Potter, 1995).

Further evidence that attention can really deal with multiple targets presented in rapid succession comes from Reeves and Sperling (1986; see also Sperling & Reeves, 1980). In these experiments, participants were presented with two concurrent RSVP streams of characters, one on the left of fixation and containing letters, the other on the right of fixation and containing digits. The task was to monitor the left stream for a cue to switch to the other stream (e.g., the letter *C*) and, as soon as it occurred, to start collecting digits from the right stream. Streams could run at various speeds (including those used in typical attentional blink studies), and by assessing both accuracy and order of report from the second stream, Reeves and Sperling determined which items in the second stream had received most attention since the onset of the cue. Across RSVP speeds, the data showed a remarkably consistent distribution as a function of time. After an initial delay (regarded as the time needed to detect the cue and switch to the second stream), chances of a digit being reported rose quickly and then deteriorated more slowly, such that most reports came from a period of 250–600 ms after cue onset. Figure 2b (bottom left panel) shows performance for one observer for streams running at 9.2 ms/item. The important finding here is that observers had relatively little trouble reporting multiple targets from this 250–600-ms period. Any existing attentional blink theory would predict that the very first digit identified in the second stream would cause an attentional blink for the following digits. There was no sign of an attentional blink. On the contrary, according to Sperling and Reeves (1980, p. 359), "the distribution directly estimates the moment of maximal attention. . . i.e., the moment of fastest registration—the moment at which the shutter is fully open." In Reeves and Sperling (1986), this performance function was described as the opening and closing of the attentional gate to short-term
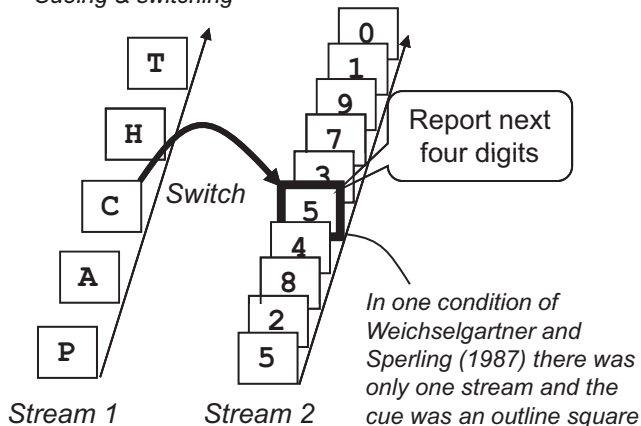
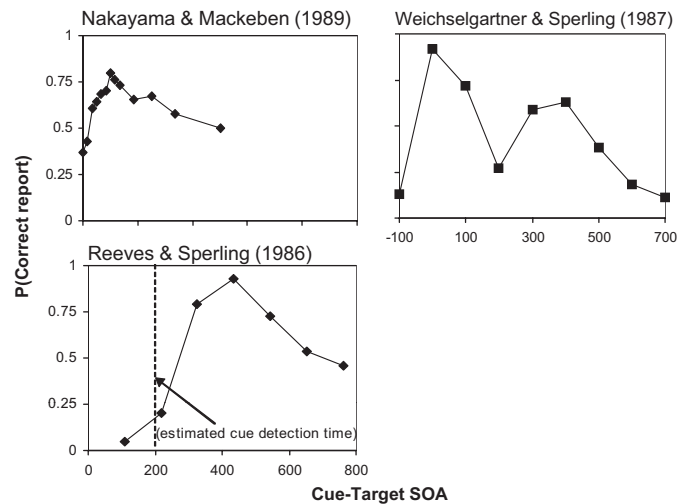## A. Typical RSVP/cueing tasks

> Nakayama & Mackeben (1989)
Cueing
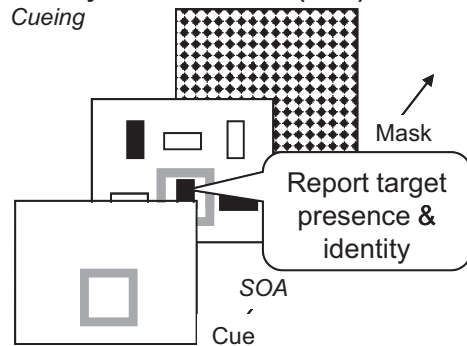
B. Typical data

> Reeves & Sperling (1986)
> Weichselgartner & Sperling (1987)
Cueing & switching

C. Typical explanation:
Transient attentional enhancement
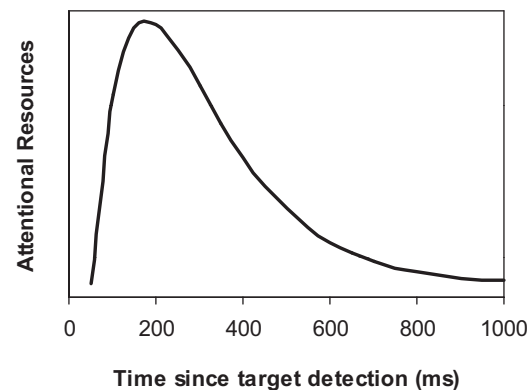
Figure 2. (A) Illustration of the tasks used by Nakayama and Mackeben (1989; top panel) and Sperling and colleagues (Reeves and Sperling, 1986; Weichselgartner & Sperling, 1987; bottom panel). All paradigms resemble a combination of rapid serial visual presentation task (RSVP; Lawrence, 1971) and cueing. In Nakayama and Mackeben, observers saw a cue, a target, and a display-wide mask. Cue and target always appeared in the same position. The task was to determine the presence and identity of one of two possible targets. In Reeves and Sperling, observers monitored one RSVP stream for a cue ("C") to switch to a second stream, from which they were required to report the next four digits they saw. In Weichselgartner and Sperling, observers monitored a single stream for an outline square (or another type of cue), after which they had to switch to reporting digits from the remainder of the stream. (B) Top left panel: Averaged accuracy data for subject NW of Nakayama and Mackeben (as estimated from their Figure 8). Top right panel: Average target report accuracy for participant EW (as estimated from Weichselgartner & Sperling's Figure 2b). Bottom left panel: Average target-report accuracy for participant AR in a dual-stream RSVP task running at 9.2 ms/item, after the observer has switched from the first to the second stream (as estimated from Reeves & Sperling's Figure 3). (C) Both transient attentional enhancement and transient attentional gating accounts hypothesize the temporary recruitment of attentional resources after a target has been detected (a gamma distribution that returns in both Reeves & Sperling and Nakayama & Mackeben, with different assumed delays). Data adapted from "Sustained and Transient Components of Focal Visual Attention," K. Nakayama & M. Mackeben, 1989, *Vision Research, 29,* Figure 8, p. 1683. Copyright 1989 by Elsevier; from "Attention Gating in Short-Term Visual Memory," A. Reeves & G. Sperling, 1986, *Psychological Review, 93,* Figure 3, p. 184. Copyright 1986 by the American Psychological Association; from "Dynamics of Automatic and Controlled Visual Attention," E. Weichselgartner & G. Sperling, 1987, *Science*, *238,* Figure 2, p. 779. Copyright 1987 by the American Association for the Advancement of Science.

memory and was modeled with a gamma distribution with a relatively steep rise (within 150 ms after cue detection) and a relatively slow decay (lasting several hundreds of milliseconds), just like the function drawn in Figure 2c. Later, Sperling and Weichselgartner (1995; see also Shih & Sperling, 2002) described it as a transition function from one attentional episode to the next.

Boost and Bounce Theory

Thus, any theory of temporal attention must deal with a major paradox: Whereas the attentional blink appears to indicate that the attentional resources available for perceptual input are reduced when a relevant event is encountered, the transient attentional-enhancement results appear to indicate that additional attentional resources are being recruited. What is more, the time courses of these two functions are remarkably similar. In fact, we have plotted the two time course functions in Figures 1c and 2c as vertical mirror images of each other. None of the current theories can simultaneously explain both these phenomena, simply because these theories start from opposite premises: the reduction versus the recruitment of attentional resources.

Our boost and bounce theory seeks to integrate the two phenomena. In the present article, we present a new theory within which this similarity in time course is no coincidence. We call this theory the boost and bounce theory of temporal attention, after its two crucial but straightforward functions: Attention boosts the visual input by responding in an excitatory manner whenever relevant information (a target) is encountered. It blocks, or bounces, the visual input in an inhibitory manner whenever irrelevant information (a distractor) is encountered. The theory claims that performance is eventually determined by the interaction between these two functions and the stimulus input. It is important to note that there is no role for capacity limitations or resource depletion in explaining the attentional blink, and its apparently long time course (in the order of 500 ms) is the result of underlying microdynamics operating at a much smaller time scale (in the order of 100 ms).

The theory shares a number of ideas with the TLC account of Di Lollo et al. (2005)—namely the importance of an input filter and a crucial role for distractors—but eliminates the necessity to invoke limited-capacity resources. The theory shares even more with the *temporary suppression account* proposed by Raymond et al. (1992), which, incidentally, was the very first theory of the attentional blink. According to the temporary suppression account, T1's defining property opens an attentional gate to higher systems so that it can be recognized. However, the post-T1 distractor is also allowed to enter, potentially leading to a false conjunction with the target-defining feature and, hence, false identification. The attentional blink occurs because this potential for conjunction errors calls for a period of suppression (or closing of the gate) to protect target processing. This then goes at the expense of later targets. For several reasons, Raymond and colleagues abandoned the suppression account in favor of a limited-capacity account (Raymond et al., 1995; Shapiro & Raymond, 1994; Shapiro, Arnell, & Raymond, 1997). First, the inhibition was thought to operate "at a relatively early stage of processing" (Raymond et al., 1992, p. 854), whereas later evidence clearly suggested that blinked items reach high levels of processing, including the activation of semantic representations (Anderson, 2005; Chua et al., 2001; Luck, Vogel, & Shapiro, 1996; Maki, Frigen, & Paulson, 1997; Marois,

Yi, & Chun, 2004; Martens, Wolters, & Van Raamsdonk, 2002; Potter, Dell'Acqua, Pesciarelli, Job, & Peressotti, 2005; Rolke, Heil, Streb, & Hennighausen, 2001; Sergent, Baillet, & Dehaene, 2005; Shapiro, Driver, Ward, & Sorensen, 1997; Shapiro, Caldwell, & Sorensen, 1997; Visser, Merikle, & Di Lollo, 2005; Vogel, Luck, & Shapiro, 1998). Second, the hypothesis that T1 needs protection from conjunction errors predicts that if no conjunction errors are possible (e.g., when targets and distractors are from different categories, or T1 involves a detection task, rather than an identification task), there should be no blink. This prediction did not hold either (Chun & Potter, 1995; Shapiro et al., 1994). Finally, there appeared to be a logical inconsistency in the idea that the post-T1 distractor triggered the inhibition: "[I]t is difficult to envision a mechanism capable of inhibiting a stimulus that itself caused the suppression" (Raymond et al., 1995, p. 661).

However, in our view, the baby has been thrown out with the bathwater. The temporary suppression model can be saved if we drop the assumptions that (a) suppression occurs only early in the system, (b) suppression is necessary to protect T1, and (c) an item cannot trigger its own suppression. In a way, then, our work serves to rehabilitate the temporal suppression account, be it in a different guise. In what follows, we explain the theory and show how it is implemented in a computational model. We then demonstrate that this model simulates a substantial number of studies on transient attention, as well as the attentional blink. We also show that the theory generates predictions that distinguish it from limited-capacity theories. Finally, we discuss the crucial differences with existing computational versions of those theories.

Stage 1: Sensory Processing

In line with Chun and Potter (1995), and many other attention theories, our theory assumes two major information-processing stages, both of which are illustrated in Figure 3. Although we refer to these as stages, they are not strictly sequential but, rather, interact. The Appendix describes a computational model in which these interactions are implemented. The first stage, which we refer to as *sensory processing*, consists of the activation of representations of perceptual features, such as color, shape, and orientation, but also of high-level representations involving semantic and categorical information. We assume that these different properties are represented separately, for example, in different feature maps that feed into one another and together form a hierarchy from simple to complex processing (e.g., Treisman & Gelade, 1980; Ungerleider & Mishkin, 1982; Zeki, 1978). We also assume that these initial sensory signals spread through the system in a mainly feedforward manner (e.g., Lamme & Roelfsema, 2000).

In line with physiological evidence, representations are activated rapidly and relatively strongly at the onset of the stimulus (the initial *visual* transient, not to be confused with transient attention here) and then decay to a more sustained level (when the stimulus remains on) or resting level (when the stimulus is switched off) of activation (e.g., Breitmeyer & Ganz, 1976). Multiple representations can be activated in parallel, but because of items being presented in the same location in RSVP, an individual item's activation is affected by preceding and succeeding items (following, e.g., Keysers, Xiao, Földiák, & Perrett, 2001; Macknik & Livingstone, 1998; Rolls & Tovee, 1994). The strength of forward masking effects on an item is assumed to depend on its similarity to the preceding item, whereas the strength of back-
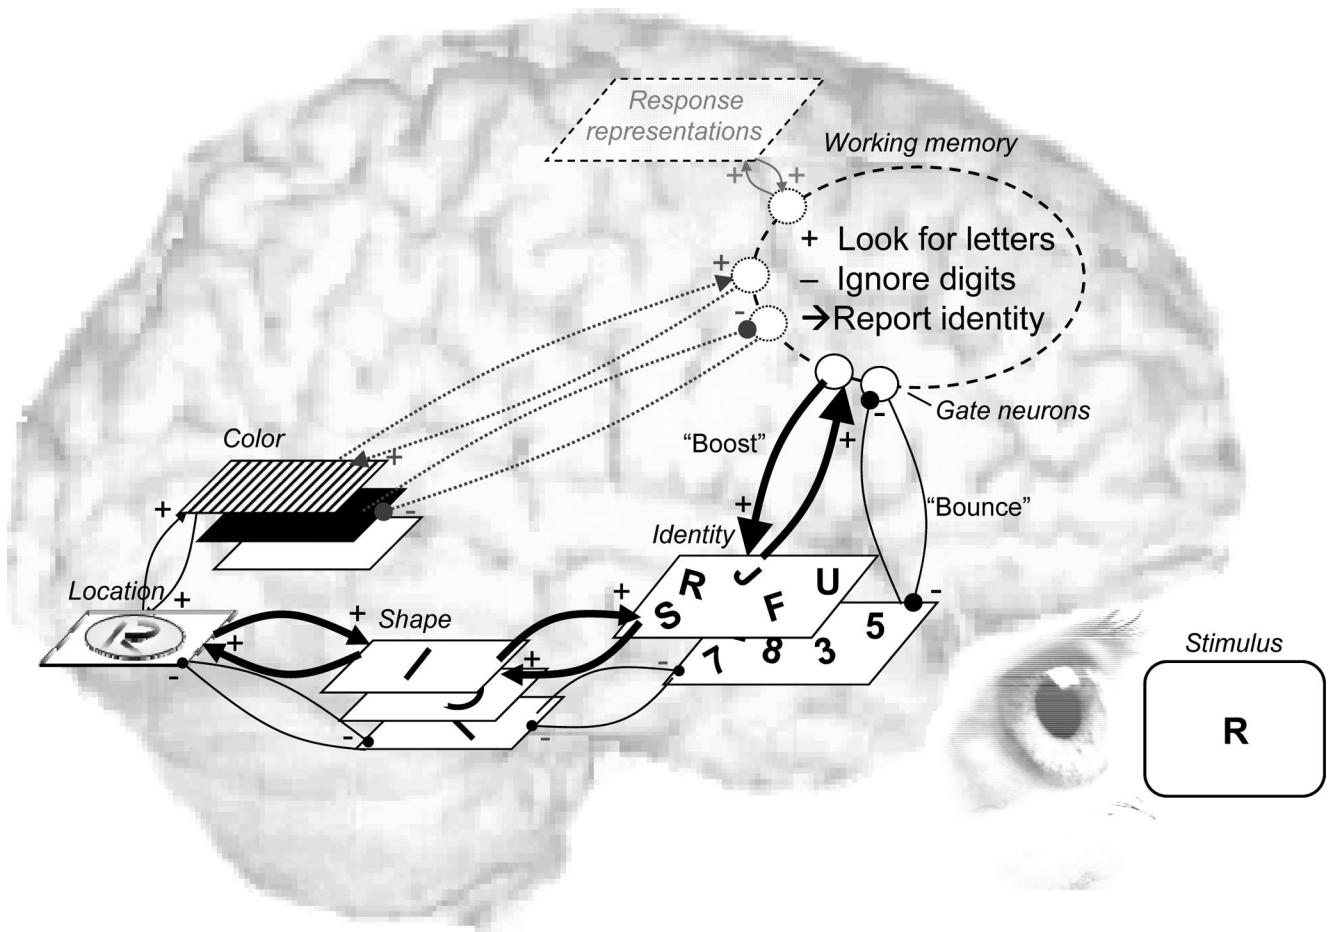
Figure 3. Diagram of the boost and bounce theory. Stimuli in the rapid serial visual presentation task (Lawrence, 1971) stream are subject to sensory processing, including activation of color, shape, and semantic properties. The attentional set required for the task is implemented in the working memory gating system, which is a combination of excitatory and inhibitory gate neurons maintaining feedback loops that respectively modulate the target- and distractor-related sensory activity. In this way, they open or close the gate to working memory, within which incoming information is then linked to reportable (e.g., verbal) representations. When a target arrives, strong attentional enhancement (i.e., excitatory feedback) is triggered, allowing the target to enter working memory. The gate remains open as long as relevant information enters. However, in the attentional blink paradigm, the bulk of the excitatory feedback hits the post-T1 distractor, which then triggers a strong inhibitory feedback response from the gate neurons in turn. An attentional blink is the consequence.

ward masking is assumed to depend on the similarity to the following item. This can be interpreted as masking being, in essence, affected by stimulus saliency: The more an item differs from its predecessor, the more salient it is, the less it is masked. Finally, in accordance with neurophysiology, the signal adapts, so that it is reduced when an item was already presented recently (e.g., Legge, 1978).

Stage 2: Working Memory

The second component is working memory. Working memory serves as the global workspace, central executive, or task monitor in which the rules applying to the task at hand are implemented and maintained (cf. Baars, 1989; Baddeley & Hitch, 1974; Bundesen, Habekost, & Kyllingsbæk, 2005; Dehaene, Kerszberg, & Changeux, 1998; Desimone & Duncan, 1995; Lavie, Hirst,

Fockert, & Viding, 2004; E. K. Miller & Cohen, 2001). Within our theory, this means that systems underlying working memory can flexibly monitor and maintain information and couple the relevant input to the relevant response. An item can only be reported when it enters working memory, because only then can it be linked to a response. Exactly how this stimulus–response mapping occurs is an important question and not one we are able to answer here. We simply assume that after receiving verbal instructions and some practice trials, the crucial links have been established. (In other words, we cannot yet fully ban the homunculus, but once he has wired the controls at the start of the experiment, he dozes off and plays no further role in our model.) In line with Dehaene et al. (1998), we envisage that these links are established through feedback connections from central working memory neurons to rele-

vant sensory representations on the one hand and to relevant response representations (whether overt or covert) on the other. An item may enter working memory—that is, it is *selected*—when sufficient (i.e., above threshold) evidence for its presence has accumulated during the time it was active in the system. The strength of this activation depends on (a) the bottom-up activity of the item, as described in the previous section, and (b) the top-down attentional modulation of this activity received through the feedback connections, as is described next.

Gating

Not only does working memory link relevant stimuli to a response but it also prevents irrelevant information from interfering with behavior by shielding response systems from unfiltered sensory representations (e.g., Engle, Conway, Tuholski, & Schisler, 1995; Fuster, 1997; Hasher & Zacks, 1988; Jonides, Smith, Marshuetz, Koeppe, & Reuter-Lorenz, 1998; Lavie, 2000). In other words, working memory employs an input filter, or attentional set, for relevant information and against irrelevant information. Others have also used the term *template matching*, but we use the term *gating* (following, e.g., Reeves & Sperling, 1986). Gating occurs through feedback that modulates sensory processing (e.g., Lamme & Roelfsema, 2000; Di Lollo, Enns, & Rensink, 2000). Sets of sensory representations that are important for the task are enhanced, or *boosted*, through excitatory feedback, whereas sets of sensory representations that are irrelevant are tempered, or *bounced*, through inhibitory feedback. Of course, the ideas of enhancing target information and suppressing distractor information are not new, and they constitute the mechanisms behind many a selective attention theory (Bundesen et al., 2005; Houghton, 1994; Lavie, 2000; Wolfe, 1994, to name but a few). We assume gating to be fuzzy. For reasons of cognitive economy, as well as flexibility, there is no one-to-one mapping between every single sensory neuron and a gate neuron (Postle, 2006). Instead, gate neurons operate on populations of sensory neurons, on relatively broad classes or categories of representations. For example, gating may be driven by "letter-like," "colored," or perhaps as specific as "reddish" stimuli but cannot apply only to something like "the little maroon espresso cup" at the exclusion of all other stimuli. In this sense, gating is also contingently *automatic*: Once an attentional set has been implemented (presumably at the start of the experiment), items that sufficiently match this set automatically elicit feedback activity, even if they are not targets (Bargh, 1992; Folk, Remington, & Johnston, 1992).

The modulation is multiplicative: It needs to interact with stimulus-based activity (Grossberg, 1995). Furthermore, feedback is issued by *gate neurons* that are themselves driven by stimuli that are predefined as targets or distractors. Enhancement and inhibition through feedback thus only really start to take off when a stimulus is present. Furthermore, we assume that gate neurons respond with an initial strong burst of activity, after which activity tapers off (e.g., Connors & Gutnick, 1990; Fuster, Bauer, & Jervey, 1982; Goldman-Rakic, 1995; Tomita, Ohbayashi, Nakahara, Hasegawa, & Miyashita, 1999; Zipser, Lamme, & Schiller, 1996). Functionally, this corresponds to the idea that updating (as opposed to maintenance) of working memory only needs to be transient (Hazy, Frank, & O'Reilly, 2006). Together with others, then, we assume that attention has a strong transient component (Nakayama & Mackeben, 1989; Weichselgartner

& Sperling, 1987; see also Bowman & Wyble, 2007, and Fragopanagos et al., 2005).

Figure 3 illustrates the boosting and bouncing feedback loops underlying the gating mechanism. When it is one's task to report letters but ignore digits, letter representations receive excitatory modulation, whereas digit representations are subject to inhibitory modulation. In both cases, although the relevant links are established, a concrete input from either category is necessary to elicit feedback from a gate neuron. As mentioned, the feedforward sweep of activity caused by stimulus presentation is assumed to travel upward through a hierarchy of representations. Similarly, top-down feedback activity travels in the opposite direction. Although initially only the alphanumeric identity representations that trigger the feedback are modulated by either boosting or bouncing feedback, it is assumed that the feedback activity rapidly spreads down the hierarchy to underlying layers of representation that originally fed into the identity representation. At the bottom of the hierarchy, this includes strong attentional modulation of the location of the feedback-eliciting item. This means that modulation is not limited to only the triggering stimulus but also affects successive items presented at the same location.

The strength of the modulation is affected by the strength of the sensory evidence. The stronger the evidence for a target or a distractor, the stronger the respective excitatory or inhibitory feedback. For this purpose, within the model, the transient attentional response is weighted by the sensory evidence accumulated during the first 15 ms of presentation. The sensory evidence itself is affected by two factors. First, the current state of attention modulates the strength of the sensory evidence, as bottom-up and top-down mechanisms continuously interact: Current excitatory feedback enhances the sensory evidence, whereas inhibitory feedback reduces it. Second, as explained earlier, the less similar targets and distractors are, the stronger the relative sensory evidence is (because of reduced masking and increased salience). This also means that when overall target salience is very high, the need for active gating is strongly reduced (although not absent), because the bottom-up activity itself provides strong evidence for what is a target and what is not. Hence, the strength of the gating is further weighted by the overall similarity between targets and distractors within the stream.

Last but not least, the spreading of feedback activity through the visual system takes time. Within our model, after sufficient activity has reached the layers of representation defining the target (i.e., a target has been detected), it takes 25 ms for it to start reaching the lowest layers. The peak of the feedback activity is not until another 70 ms have passed, and hence, the bulk of attention does not arrive until roughly 100 ms after target detection. This delay means that, in RSVP paradigms, the triggering stimulus may actually already be gone and replaced by the next stimulus by the time the full chain of recurrent processing is operational. Thus, although the goal of modulation by feedback may be to retrieve the precise identity of a target stimulus, the ensuing feedback may accidentally also, or even predominantly, modulate the trailing stimuli.

Capacity Limitations

There are two capacity limitations in our model:

1. No more than one attentional set (or task set) can be active at the same time. This entails that a reconfiguration

of selection settings, or *switch*, is necessary when there are multiple tasks, multiple locations, and possibly multiple modalities involved. New gate neurons need to be recruited or configured to filter for a different location or a different kind of stimulus. The exact mechanisms of such switches are not implemented in the model yet. For the time being, we assume that gate neurons activate one another and that this causes a transition from one attentional set to the next. Following Sperling and Weichselgartner (1995), we modeled this as follows: After a stochastically determined interval, a new set of attentional gates opens, and the old gates close. We chose a logistically distributed stochastic interval with a mean that depends on the type of switch. Endogenously cued switches (including task switches) have a mean switch time of 200 ms, whereas exogenously cued switches are assumed to take only 75 ms (Carlson, Hogendoorn, & Verstraten, 2006; Cheal & Lyon, 1991; Eriksen, 1990; Jonides, 1981; Luck & Vecera, 2002; Posner, 1980; Posner & Cohen, 1984; Shulman, Remington, & McLean, 1979).

2. Working memory storage capacity is limited. The likelihood of stimuli entering working memory is not constant throughout an RSVP stream: When capacity has been exhausted by previous items that have entered working memory, new items cannot enter. Capacity was set at five items (following the evidence that working memory capacity is about four to seven items, depending on additional memory strategies, such as chunking, rehearsal, or efficient use of additional memory systems, Cowan, 2001; G. A. Miller, 1956).

It is important to note that, within our theory, neither of these capacity limitations provides an explanation for the attentional blink, as the attentional blink easily occurs with only two targets, for which one and the same attentional set applies (e.g., Chun & Potter, 1995).

Explaining Empirical Findings: Transient Attentional Enhancement

In this and the following section, we describe how boost and bounce theory explains performance enhancement under some circumstances but performance deterioration under others. It also explains many related findings. For each finding, we demonstrate the viability of our theory through simulations with the computational model specified in the Appendix. It is important to note that these simulations were generated with just one set of parameter values (except for parameters governing similarity between items, where mentioned). Although this precludes optimizing the fit of particular individual experiments, we wished to capture the overall pattern of findings with as few assumptions as possible. The simulations should therefore be evaluated on the basis of their qualitative, not quantitative, fit of the data pattern.
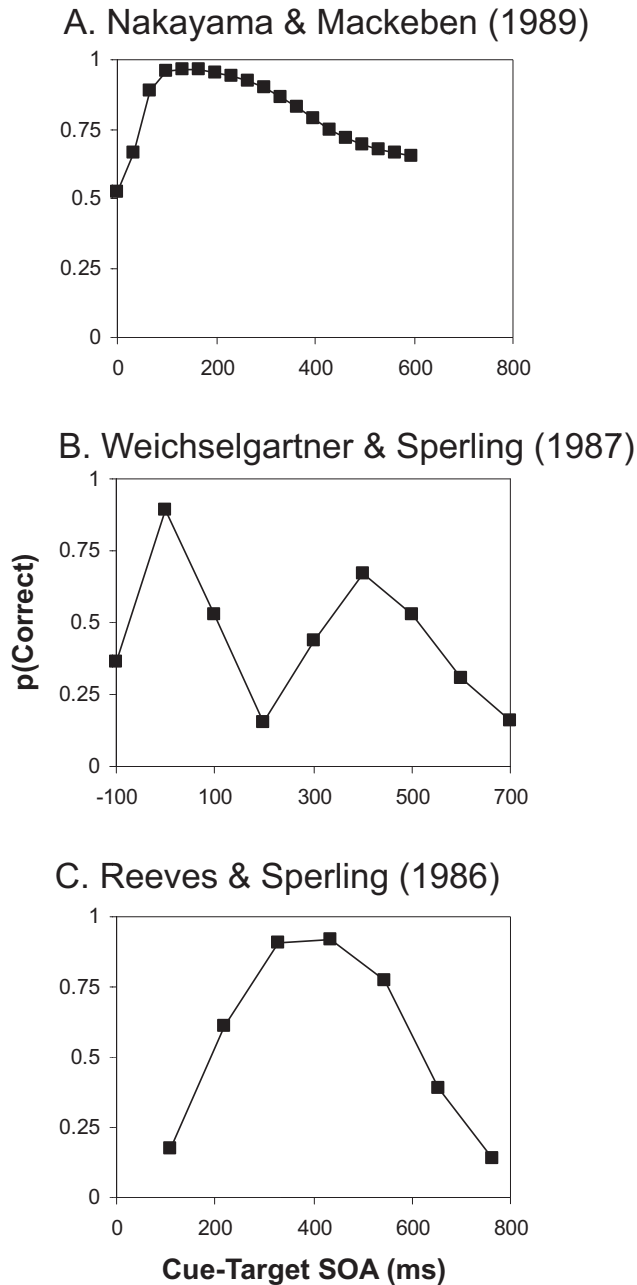
Transient Spatial Cueing Effects

A relevant event (usually a target or a valid cue) elicits excitatory feedback from gate neurons to the sensory neurons activated by the event. This results in a transient boost of the neuronal signals coding for this event, including representations of its features and source location. When the gate neurons are triggered, the enhancement is initially strong but then decays back to normal levels. This mechanism explains the results of the cueing study of Nakayama and Mackeben (1989), in which observers had to detect a target in a cluttered and masked display. A 100% valid abrupt onset cue indicated the target position. Performance first rose rapidly with increasing time between cue and target, but for cue–target intervals greater than about 100 ms, it slowly decreased again, while remaining above chance levels. Because in Nakayama and Mackeben's study, the target and mask were always in the same position as the cue, the largest part of their task can be simulated as a straightforward RSVP. However, because participants did not know where the abrupt onset cue would appear, we assume that the initial part of processing was dominated by an exogenously triggered spatial switch necessary to orient to the new location. Within our model, such a switch takes on average 75 ms. Figure 4a shows the performance predicted by the model for a target presentation duration of 33 ms (which is the duration Nakayama and Mackeben used for most of their participants).[1] Because the cue remained on during target presentation, and the target was presented within a cluttered array and followed by a strong mask, we assumed masking to be stronger here (by increasing the similarity by a factor 5) than in the subsequent RSVP paradigms.

Two Glimpses

The idea of a rapid transient attentional enhancement had surfaced earlier, in a study by Weichselgartner and Sperling (1987), who referred to it as the first of two attentional "glimpses." The second of the two glimpses was thought to reflect a slower, more sustained attentional process independent of the first glimpse. In the crucial condition, observers monitored a single stream of digits for a digit that was uniquely accompanied by a salient cue (e.g., a surrounding square or a highlighted numeral). The task was to report this cued digit, as well as the digits that followed it. The pattern of performance showed two peaks: There was relatively good performance for the initial target and often the item presented immediately after, as well as for items presented after about 300–400 ms (after which performance gradually decreased). Between these two peaks, performance showed a brief but marked drop. As shown in Figure 4b, our model reproduces these findings. However, within our model, the "two glimpses" are not the result of two independent attentional processes being separated in time but of the continuous interaction of the boosting and bouncing gating functions. Note that in Weichselgartner and Sperling's task, observers needed to ignore digits until the cue (the surrounding square) was presented, after which digits were to be reported. This has two implications in our model: (a) Gate neurons are initially

---

[1] Although the Nakayama and Mackeben task involved a spatial cue, it is mainly informative about the time course of attention, rather than spatial selection. This is because the target (and, in a sense, the display-wide mask) always appeared in the same location as the cue. It is interesting to note that the time course of transient attention does show remarkable similarities to that of inhibition of return in more classic spatial cueing paradigms (Maylor & Hockey, 1985; Posner & Cohen, 1984). Whether or not transient attention is an important factor in inhibition of return remains an important question for the future.

A. Nakayama & Mackeben (1989)

B. Weichselgartner & Sperling (1987)

C. Reeves & Sperling (1986)

Cue-Target SOA (ms)

Figure 4. Model simulations regarding the attentional enhancement findings of (A) Nakayama and Mackeben (1989), (B) Weichselgartner and Sperling (1987), and (C) Reeves and Sperling (1986). In Simulation A, backward masking of the target was assumed to be stronger (by a factor of 5) than in rapid serial visual presentation paradigms, as Nakayama and Mackeben used cluttered target displays with relatively strong, long duration masks. In Simulation B, the cue (outline square, highlighted numeral) was assumed to be more salient than typical rapid serial visual presentation items, and its color similarity was therefore reduced to 0.5. SOA = stimulus onset asynchrony.

set to filter for the additional feature provided by the cue (the outline square). This means that, initially, items with cue are treated as targets and cause enhancement, whereas items without cue are treated as distractors and cause inhibition. (b) After the

cue, digits are to be treated as targets, and an endogenous shift in attentional set toward accepting digits is now required. Such an endogenous task switch takes time. The dynamics of the model are then as follows: The presentation of the cue results in a boost of the identity of the simultaneously presented digit, which enters working memory. Because of inherent delays, the boost reaches its maximum level when the item *after* the cued item is presented. Hence, the identity of the post-cue item often also enters working memory, despite the fact that, within the initial attentional set, it is still a distractor. Together, these first items in the target series constitute the first glimpse. At the same time, the cue also acts as a signal to initiate the task switch, which takes 200 ms on average (because it is endogenous). During this period, gate neurons engaged with the first task set have not yet switched off, and the boosted post-cue item sets off inhibition exactly because it is still being treated as a distractor. In other words, the attentional gate closes for the next item, resulting in a dip in performance. Later on, we see that this inhibition, as set off by the post-target distractor, plays an important role in explaining the attentional blink. In fact, the brief performance dip in the Weichselgartner and Sperling results can be regarded as an attentional blink beginning to develop (cf. Raymond et al., 1992) but then being rapidly reversed as the second task is activated. This is because, as time passes, the attentional set is more and more likely to have switched to one in which digits without accompanying cues are treated as targets. The attentional gate reopens, and performance rapidly recovers. Excitatory feedback, although already subsiding, is still active and catches some of these targets in its tail: A second glimpse is observed (see Figure 4b). Finally, as working memory fills up with targets, new items are increasingly unlikely to enter working memory. The model's prediction then is that, when task switches are removed, only a single glimpse should emerge. Indeed, recent data (Olivers, Van der Stigchel, & Hulleman, 2007; Di Lollo et al., 2005) appear to support this prediction.

Switching Streams

Figure 4c shows how the model accounts for the item-report data from observer AR in Reeves and Sperling (1986) for an RSVP task consisting of two streams running at 9.2 items/s. In this task, observers were required to switch streams as well as tasks after a cue in the first stream (the cue itself was not to be reported). The data show the proportion of correctly reported items for each temporal position in the second stream, after the cue. Within the model, the first 200 ms of performance are therefore dominated by the fact that working memory needs to switch from filtering for cue-defining properties to filtering for digits in the second stream. Unlike the Weichselgartner and Sperling (1987) task, in this task, no first glimpse is observed, because as long as the attentional switch has not occurred, participants are still attending the first stream, which does not contain any reportable items. Then, as the switch to the second stream is completed, the targets start to enter and fill up working memory. Because there are targets in only the second stream, only excitatory feedback is operational, and there is no inhibition. Nevertheless, performance eventually deteriorates simply because working memory is full.

Explaining Empirical Findings: The Attentional Blink

We consider the typical attentional blink task of having to report two letters from a stream of digits, as it involves no spatial or task switches. Excitatory feedback links are set up for letters, and inhibitory links for digits, although note again that these categories are considered to be fuzzy. The stream is running, but T1 has not been presented yet: Only distractors are encountered, and during the first few, the system settles in a stable inhibitory state sufficient to keep them out of working memory. Direct evidence for such a pre-T1 inhibitory state comes from Dux, Coltheart, and Harris (2006); Maki and Padmanabhan (1994); Olivers and Watson (2006); and Sahraie, Milders, and Niedeggen (2001).

Then, after the initial streak of distractors, T1 appears. Sensory evidence for a target arrives in the neurons coding for targets, which, in turn, elicits feedback from the gate neurons in working memory linked to these neurons. The result is a surge of excitatory recurrent processing, boosting sensory signals for T1 and allowing T1's identity to enter working memory. Figure 5 illustrates the dynamics of attention following T1. The most important aspect of the model is that, even though T1 benefits, the peak of the boost in excitatory feedback actually arrives *after* T1, especially at the lower layers of the representational hierarchy, which need some more time before they are reached by the recurrent wave of activity. At the usual presentation rates of 100 ms/item, this means that the bulk of attention lands on the post-T1 item.

The attentional blink occurs when this post-T1 item is a distractor. Because of the delay in excitatory feedback, this distractor accidentally receives the maximal enhancement intended for T1. As a consequence, the gate neurons now receive a very strong signal of the wrong category entering working memory. Because the gating system has been set up to reject distractors, this strong signal automatically results in a *bounce*—a strong inhibitory response, which, in effect, closes off working memory from the RSVP stream. As the inhibition builds up over the course of the post-T1 distractor's presentation, the distractor itself still has a reasonable chance of entering working memory, resulting in substitution errors when the task allows for such errors to be made (cf. Isaak et al., 1999). However, just as the boost reaches its maximum after T1, the bounce reaches its maximum after the post-T1 distractor. If this next item turns out to be a target, an attentional blink arises: Because of the strong inhibitory feedback, this target is less likely to muster sufficient activity in time to make it into working
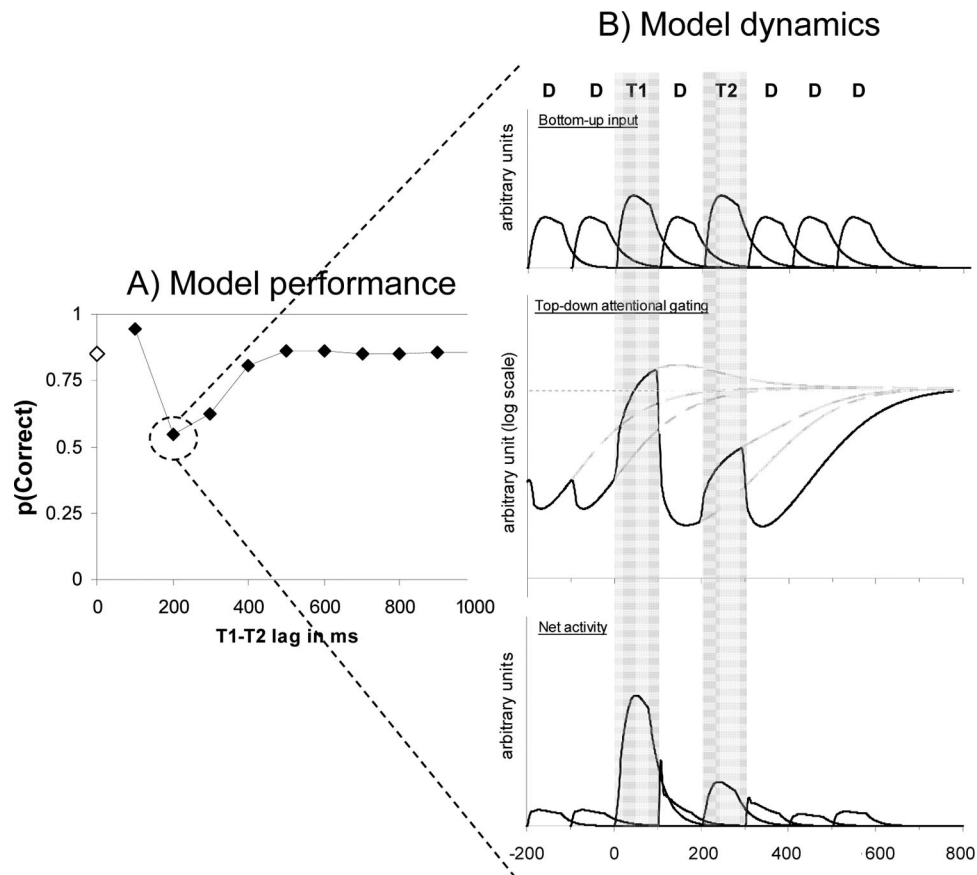
Figure 5. Simulation of the basic attentional blink paradigm. (A) The model's performance is shown for T1 (the open marker on the *y* axis) and for T2 as a function of T1-T2 lag. (B) The model's dynamics for a T2 at lag 2 is shown in detail, with (1) the bottom-up sensory signal, (2) the top-down attentional response to the input (as a combination of both excitatory and inhibitory feedback), and (3) the combined bottom-up and top-down signal.

memory. This is the case even though this target itself generates some net attentional boost (see Figure 5). Over time, as the strong transient inhibition summoned by the post-T1 distractor gradually resides, performance gradually improves.

Thus, the attentional blink is completely determined by the temporal dynamics of excitatory and inhibitory selection mechanisms following T1 and its trailing distractor(s). It is not determined by the need to process and consolidate T1 within working memory, nor is it determined by limited working memory capacity. Because the most important factor is the delay in feedback processing, one prediction is that the attentional blink is mainly time-based, not item-based (even though it is not entirely time-based, because distractor items are still necessary to trigger the inhibition). If, for example, the stream runs at twice the speed, twice as many items should fall within the blink. This prediction has recently been confirmed (Bowman & Wyble, 2007; Martens, Munneke, Smid, & Johnson, 2006; Popple & Levi, 2007; see Nieuwenhuis, Gilzenrat, Holmes, & Cohen, 2005, for a partial manipulation), and, as shown in Figure 6, our model reproduces these findings.

The past decades have generated a wealth of other findings related to the attentional blink, many of which we now discuss.

Lag-1 Sparing

Probably the most important phenomenon accompanying the attentional blink is lag-1 sparing. Performance for T2 is typically high when T2 immediately follows T1 (e.g., Chun & Potter, 1995; Raymond et al., 1992; see Figure 1). In fact, at lag 1, T2 performance often exceeds T1 performance (Chun & Potter, 1995; Di Lollo et al., 2005; Olivers et al., 2007; Potter, Staub, & O'Conner, 2002). This has been a troublesome finding for most limited-capacity theories. How can T2 be spared when only one item can pass through the bottleneck, or when resource depletion should be maximal? So far, solutions have been sought in additional hypotheses and psychological constructs that allow for the two targets to be temporarily processed together (Chun & Potter, 1995), bound to the same token (Bowman & Wyble, 2007; Chun, 1997b) before a sluggish attentional gate closes (Visser, Zuvic, Bischof, & Di Lollo, 1999), in a single batch (Jolicoeur et al., 2002), temporal window (Visser, Bischof, & Di Lollo, 1999), glimpse

(Chua et al., 2001), event representation (Kessler et al., 2005), or episode (Hommel & Akyürek, 2005). We do not deny the existence of such temporary attentional processes—in fact, many are very similar to the temporary enhancement function within our model. The point here is that within most of these theories, the hypotheses of what causes lag-1 sparing do not explain the attentional blink itself, and vice versa. As Visser, Bischof, and Di Lollo (1999) wrote, the attentional blink and lag-1 sparing are treated as independent phenomena: T1 triggers a relatively long episode of up to 500 ms during which resources are taken away from the stream, resulting in the blink. At the same time, T1 also triggers a shorter period of about 200 ms, during which the lag-1 item is assigned additional resources.

In contrast, our theory provides an integrative explanation of both the attentional blink and lag-1 sparing, by proposing the T1-induced attentional boost as the common cause. The post-T1 item is spared and may even yield better performance than T1 itself, simply because it arrives in the peak of the attentional boost. Because of the boost's set delay, the theory predicts that lag-1 sparing is also largely time based, rather than item based. In other words, relative sparing should not be confined to lag 1, as long as the later lags fall near the peak of the attentional boost. The few studies that have looked at the issue all suggest that the sparing is indeed essentially time based and spreads to later lags when the RSVP is speeded up (Bowman & Wyble, 2007; Martens, Munneke, et al., 2006; Nieuwenhuis et al., 2005). Again, Figure 6 shows that this pattern of results is successfully reproduced by our model.

Apparent Tradeoffs Between T1 and T2

Potter et al. (2002) systematically varied the SOA between T1 and T2 and found evidence consistent with the idea that T1 and T2 are in direct competition for resources. At short SOAs (<100 ms), T2 performance was better than T1 performance (a result indicative of sparing), whereas at long SOAs (>100 ms), T2 performance was inferior to T1 performance (as is indicative of an attentional blink; see also Bachmann & Hommuk, 2005). However, in this study, SOA was confounded with differential masking effects. The targets were words, whereas the distractors surrounding the targets were always ampersands. Thus, at very short SOAs, T1 was probably much more effectively masked by the closely following and very similar T2 than
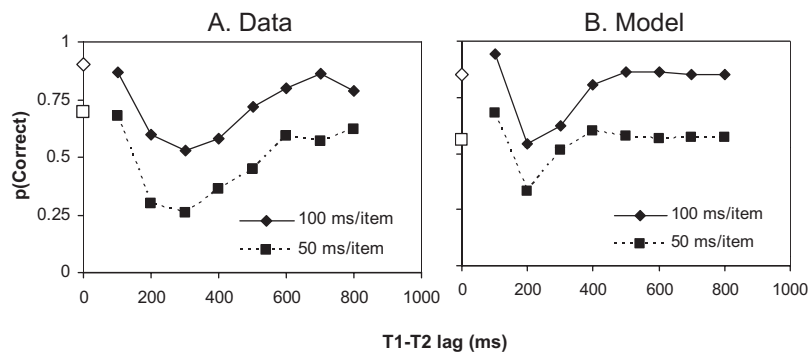
Figure 6. T2 performance as a function of T1-T2 lag in streams with interitem stimulus onset asynchronies (SOAs) of 50 ms, versus of 100 ms. T1 performance is shown by open markers on the *y* axis (lag 0). (A) Data adapted from "The Simultaneous Type, Serial Token Model of Temporal Attention and Working Memory," H. Bowman & B. P. Wyble, 2007, *Psychological Review, 114,* Figure 19, p. 55. Copyright 2007 by the American Psychological Association. (B) Model results.

at long SOAs, when it was followed by a train of dissimilar distractors, whose representation was probably further weakened because of their repetition (Kanwisher, 1987). T2, on the other hand, suffers more from the attentional blink at later lags, when more distractors intervene. Such differential effects emerge from our model, and, as can be seen in Figure 7, it captures the large part of the apparent tradeoff between T1 and T2, without the necessity to assume a real tradeoff in terms of attentional resources (see also Kawahara & Enns, in press).

Spreading the Sparing

We have already mentioned that under sufficiently rapid presentation rates, sparing can spread to lag 2. But even at standard presentation rates of around 100 ms/item, sparing is not necessarily limited to lag 1. Di Lollo et al. (2005), as well as Olivers et al. (2007), asked participants to identify the targets in sequential triplets of items, like . . . *TDT*. . ., and . . . *TTT*. . . (*T* denoting a target, *D* denoting a distractor; all embedded in a stream of distractors). Note that the final targets in these triplets are in exactly the same temporal position relative to the first target, and a limited-capacity account would thus predict an attentional blink for both these final targets. Yet performance differed remarkably: There was a clear blink for the final target in the *TDT* triplet, whereas there was no blink for any of the targets in the *TTT* triplet (this relative sparing of the third target even occurred when performance was analyzed contingent upon T1 and T2 correct). In other words, lag-2 sparing occurred. Olivers et al. (2007) showed that with a total of four targets, sparing spreads even further, to lag 3. Work by Nieuwenstein and Potter (2006) demonstrated that sparing may spread to five subsequent targets. These results are highly problematic for resource-depletion theories. They predict that performance should deteriorate, not improve, with increasing numbers of targets. Boost and bounce theory, on the other hand, readily explains these findings: As long as no distractor signal arises, no inhibition is issued (the gate remains open), allowing for multiple target sparing.

Figure 8 shows the lag-2 and lag-3 sparing found by Olivers et al. (2007) and the model's replication.

Rapid Reversal of the Attentional Blink

Even more problematic for limited-capacity accounts is the finding that sparing can still occur once a proper attentional blink has been induced. In other words, the attentional blink can be rapidly reversed. Olivers et al. (2007; see also Kawahara, Kumada, & Di Lollo, 2006) presented observers with the sequence . . . *TDTT*. . . . As expected, an attentional blink was found for the second of the three targets. However, the third target was almost completely spared, even though it was presented in a temporal position relative to the first target that would normally be severely "blinked" (as was confirmed using . . . *TDDT*. . . sequences). Similar findings have been reported by Nieuwenstein, Chun, Van der Lubbe, and Hooge (2005) and Nieuwenstein (2006). Just prior to T2, they inserted distractors carrying the target-defining property (referred to as *cues*, these distractors could, e.g., be red when targets were also red). Again, substantial sparing occurred.

The reversal of the blink also appears problematic for the TLC account (Di Lollo et al., 2005). TLC hypothesizes that, during the attentional blink, the central executive control is occupied by T1, which leaves the input filter vulnerable to exogenous disruption. Within its current form, TLC does not allow central control to be reinstated before T1 processing has finished. Instead, T2 might cause an exogenous reset of the filter to start accepting targets again (as is proposed by Kawahara, Kumada, & Di Lollo, 2006), although it is difficult to see how these targets are recognized as such when they are not allowed to pass the filter in the first place. Moreover, Nieuwenstein (2006) reported strong evidence that the attentional filter is actually fully intact during the attentional blink. He found that when observers searched for red targets in an RSVP stream, T2 was relatively spared when it was preceded by a red cue but was hardly spared when it was preceded by a green cue. However, when he asked observers to detect both red and green targets, both red and green cues became effective, regardless of the actual color of T2 itself. TLC cannot account for this, because green cues should have exogenously reconfigured the input filter to accept only green targets, not red targets.

The above results imply that, rather than reflecting the irreversible T1-induced draining of resources for 500 ms, the attentional blink reflects the dynamic, relatively rapid, online response to important changes in the stimulus stream. Given the speed at which these streams run (typically at 10 items/s), the data suggest that the system responds on a time scale of something more like 100 ms, rather than 500 ms. Boost and bounce theory readily allows for such rapid reversals of the attentional blink following either targets or cues. A particular target (or target-like cue) is
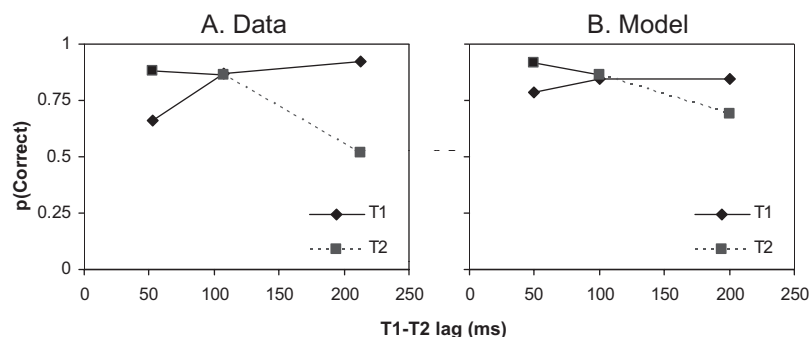

Figure 7. (A) T1 and T2 identification performance, as a function of T1-T2 lag, adapted from "The Time Course of Competition for Attention: Attention Is Initially Labile," M. C. Potter, A. Staub, & D. H. O'Conner, 2002, *Journal of Experimental Psychology: Human Perception and Performance, 28,* Figure 1, p. 1152. Copyright 2002 by the American Psychological Association. (B) Model results.
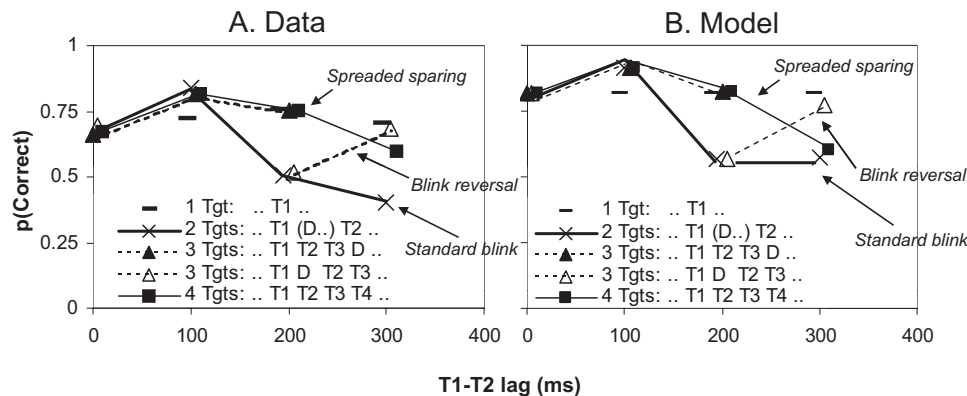
Figure 8. (A) Target identification performance in an attentional blink paradigm in which either one, two, three, or four targets were presented, as adapted from "Spreading the Sparing: Against a Limited-Capacity Account of the Attentional Blink," C. N. L. Olivers, S. Van der Stigchel, & J. Hulleman, 2007, *Psychological Research, 71,* Figure 3, p. 132. Copyright 2007 by Springer-Verlag. Accuracy is shown for each possible temporal position, with T1 appearing at temporal position 0; T2 appearing at temporal positions 1, 2, or 3 (at 100, 200, or 300 ms, respectively); T3 at temporal positions 2 or 3 (200 or 300 ms, respective); and T4 at temporal position 3 (300 ms). The data show that sparing spreads to lags 2 and 3 and that the blink can be rapidly (i.e., within a single lag) reversed. (B) Model results.

blinked because the system is still in a strong inhibitory state following a distractor. Nevertheless, blinked targets or cues do not go completely unnoticed by the system and can already start eliciting excitatory feedback. The next item, when it is another target, benefits from this, resulting in an increased sparing. If the next item turns out to be a distractor instead, substitution errors may occur (Chun, 1997a; Isaak et al., 1999). As shown in Figure 8, the blink's reversal is indeed found in our simulation of the Olivers et al. (2007) paradigm. In this simulation, we set $d_{catsame}$ equal to d_{catdif} (0.7) to model the fact that Olivers et al. had equalized masking across targets and distractors. As shown in Figure 9, the cueing data of Nieuwenstein et al. (2005) are, to a large extent, also captured. Here, it is assumed that a distractor carrying the target-

defining property is treated as neither a distractor nor a target but as an item with neutral relevance.

The Importance of the Distractors

As mentioned earlier, to explain lag-1 sparing, some have suggested that whatever happens to the lag-1 item is independent of what causes the attentional blink. However, this is directly at odds with the evidence that the lag-1 item, when it is a distractor, is of crucial importance in inducing an attentional blink in the first place. In the previous section, we discussed the finding that . . . *TDT* . . . sequences generate a blink for the final target, whereas . . . *TTT* . . . sequences do not (Di Lollo et al., 2005; Olivers et al.,
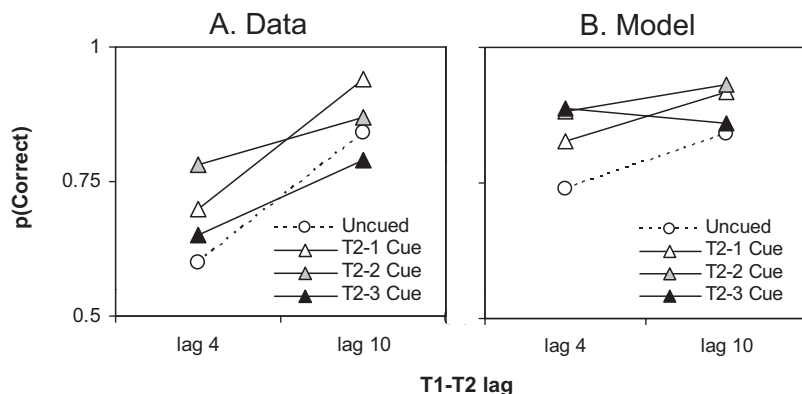
Figure 9. (A) T2 identification performance for red targets among black distractors as a function of T1-T2 lag and as a function of the number of preceding distractor items with the same color as the targets (i.e., the number of red cues), as adapted from "Delayed attentional engagement in the attentional blink," M. R. Nieuwenstein, M. M. Chun, R. H. J. v. d. Lubbe, & I. T. C. Hooge, 2005, *Journal of Experimental Psychology: Human Perception and Performance, 31,* Figure 2, p. 1466. Copyright 2005 by the American Psychological Association. (B) Model results.

2007). The only difference between these conditions is the presence of a distractor between the targets. Furthermore, the attentional blink is strongly reduced when the distractor following T1 is replaced with a blank (Breitmeyer, Ehrenstein, Pritchard, Hiscock, & Crisan, 1999; Chun & Potter, 1995; Grandison et al., 1997; Raymond et al., 1992; Seiffert & Di Lollo, 1997), or when the distractor following T1 is weakened because of inhibition or repetition blindness (Drew & Shapiro, 2006; Dux et al., 2006). There are also indications that the post-T1 distractor receives relatively high levels of processing compared with other items in the stream. For example, Chua et al. (2001) found that only distractors presented at lag 1 could semantically prime a subsequent T2. As mentioned earlier, instead of T1, observers often report the immediately following distractor (provided the type of task allows for such errors to be made; Chun, 1997a; Isaak et al., 1999; Raymond et al., 1992). At the same time, responses to a target at the very end of the stream have been found to be inhibited when it matches the post-T1 distractor (Loach & Marí-Beffa, 2003). This suggests that the post-T1 distractor may first be enhanced and then be suppressed.

Boost and bounce theory readily predicts all these findings. Within the theory, the post-T1 distractor is essential to observing a blink exactly because it induces the inhibition that causes the blink. Removing or weakening the post-T1 distractor removes or reduces the inhibition and, thus, the blink. At the same time, the post-T1 distractor itself is so heavily boosted that it may act as a strong subconscious prime or even break through the gate to consciousness.

Figure 10 shows the predictions from our model for the case in which the lag-1 distractor is removed, compared with data from Chun and Potter (1995). The model makes the additional assumption that a gap in a continuous stream is slightly disruptive (i.e., it is a salient but irrelevant event) that triggers a weak inhibitory response (at 1/10 of the normal inhibitory response). Without this assumption, the pattern looks the same, but the distractor after the gap induces a deeper attentional blink. Figure 11 shows how the attentional blink is reduced when the post-T1 distractor is weakened (Drew & Shapiro, 2006; Dux et al., 2006). Figure 5 shows how the post-T1 distractor is first enhanced and then suppressed but, nevertheless, generates substantial activity, compared with other distractors.[2]

Finally, the post-T2 distractor also plays an important role in the model. When the post-T2 distractor is removed, sensory activation associated with T2's identity lingers (as it is no longer masked). Excitatory feedback elicited by T2 reaches its maximum while this sensory activity is still around, whereas normally, it enhances the post-T2 distractor more than T2. T2 is therefore perceived at a delay, consistent with behavioral and neurophysiological data (Dell'Acqua, Pascali, Jolicoeur, & Sessa, 2003; Giesbrecht & Di Lollo, 1998; Jolicoeur & Dell'Acqua, 1998; Vogel & Luck, 2002; Zuvic, Visser, & Di Lollo, 2000).

Similarity Effects

Not only the presence but also the nature of the distractors is important. A number of studies have shown an aggravated attentional blink when either the pre-target distractors, the post-target distractors, or all the distractors in the stream, are similar to T1, T2, or both (Chun & Potter, 1995; Dux & Coltheart, 2005; Ghorashi, Zuvic, Visser, & Di Lollo, 2003; Giesbrecht, Bischof, & Kingstone, 2003; Isaak et al., 1999; Kawahara, Enns, & Di Lollo, 2006; Maki, Couture, Frigen, &

Lien, 1997; Maki, Bussard, Lopez, & Digby, 2003; Maki & Padmanabhan, 1994; McAuliffe & Knowlton, 2000; Olivers & Watson, 2006; Raymond et al., 1995; Visser, Bischof, & Di Lollo, 2004). Although it has been difficult to separate out the exact contributions, similarity effects appear to occur on the level of visual features (probably involving masking), as well as on conceptual/semantic levels. Furthermore, manipulations of target–distractor similarity often also involved variations in similarity within the distractor set itself. Finally, it is often difficult to assess whether similarity affects overall performance or specifically affects the attentional blink (i.e., interacts with lag). This is because potential ceiling effects allow less room for improvement at later lags than at the usually more affected shorter lags.

Within our model, increased similarity between items (within or between categories) reduces the sensory input signal to gate neurons and, thus, both the strength of excitatory and inhibitory feedback elicited by targets and distractors, respectively. Furthermore, as targets and distractors become less distinguishable, there is an increased necessity for top-down gating to extract the relevant items from the stream. The net result is a deeper attentional blink. Figure 12 shows simulated attentional blinks for a range of within- and between-category similarities (parameters $d_{catsame}$ and d_{catdif}, respectively).

A New Prediction: The Attentional Blink Can Be Postponed

Boost and bounce theory also makes new predictions. One prediction is that the attentional blink is not time-locked to T1 but to the first post-T1 item that does not resemble a target (and that therefore induces strong inhibition). The prediction then is that inserting post-T1 distractors that do not induce strong inhibition will lead to a postponement of the attentional blink until a distractor is encountered that does. We did this by making the post-T1 distractors carry the target-defining feature, as is illustrated in Figure 13a. We started from an attentional blink task in which observers were asked to report the two red letters in a stream of black digits. We assume that observers set up their gating system (or attentional set) such that red items led to excitatory feedback, whereas black items led to inhibitory feedback. In the *standard* condition, the red T1 was usually followed by a black distractor (except at lag 1), and a standard attentional blink pattern was expected. The interesting case is the *T1 + 1 red* condition, in which the distractor immediately following T1 was also red. Our model predicts that, simply because it carries a target feature, this distractor will not induce the inhibition invoked by the standard black distractor. In fact, the target resembling distractor is treated as neutral (neither target nor distractor) by the model, and the excitatory feedback invoked by T1 continues unhindered. The

---

[2] It is interesting to note that Seiffert and Di Lollo (1997) found a clear attentional blink, even when the lag-1 distractor was removed, as long as T1 itself was degraded by a simultaneous mask. This appears at odds with our model, which assigns a central role to the post-T1 distractor. However, if we assume that degrading T1 results in slower processing of T1, then the triggering of excitatory feedback is delayed too. In other words, instead of the lag-1 item, the lag-2 item now receives the bulk of the processing. If this item is a distractor, it will result in a blink for subsequent items. Thus, a full but delayed attentional blink is predicted, as was exactly found by Seiffert and Di Lollo (1997).
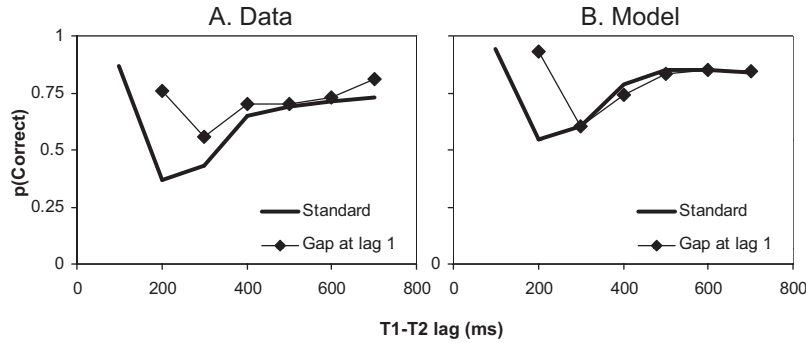
Figure 10. (A) T2 identification performance as a function of T1-T2 lag either in a standard attentional blink task or when a gap is inserted at lag 1 (in other words, the lag 1 distractor is removed), as adapted from "A Two-Stage Model for Multiple Detection in Rapid Serial Visual Presentation," M. M. Chun & M. C. Potter, 1995, *Journal of Experimental Psychology: Human Perception and Performance, 21,* Figure 4, p. 114. Copyright 1995 by the American Psychological Association. (B) Model results.

model thus predicts that the attentional blink will not be fully induced until the next black item appears, which, in this case, is the T1 + 2 distractor. In other words, the entire attentional blink function should shift later in time by one lag. Similarly, in the *T1 + 2 red* condition, we inserted two red distractors after T1. The model now predicts that the same attentional blink function should shift even further back. In sum, the attentional blink is postponed. Figure 13b shows this prediction.

Limited-capacity theories predict something rather different. First of all, no postponement is predicted, simply because the attentional blink is time-locked to T1. After all, it is the second-stage processing of T1 that is assumed to be the major culprit behind the attentional blink. Second, such theories have often assumed that T1 induces a blink because it is masked by a distractor. If anything, then, making the post-T1 distractor similar to T1 (by giving it the same color) should make it more difficult to maintain the same level of T1 performance, resulting in a deeper blink for T2. In a similar vein, Shapiro and Raymond's (1994) interference theory, with its emphasis on similarity, would predict that T2 suffers more from the red distractors than from the black distractors, simply because the red distractors match the target

template better and are therefore stronger competitors for selection. So again, a deeper blink, but no postponement, is predicted. Finally, similar to the boost and bounce model, TLC theory might predict that the input filter remains intact as long as items carrying the target feature enter the system, and an attentional blink is only induced when the filter is disrupted. However, unlike our model, TLC predicts the attentional blink to remain time-locked to T1, because it assumes that it is T1 processing that prevents the central executive from maintaining the input filter. As soon as the central executive is available again, the input filter should be reinstantiated, regardless of the type of intervening distractors.

Method

Participants. Eighteen participants (9 male) were included in the study, ranging in age from 18 to 33 years (*M* = 22.3 years). They were paid €7 (about $11) per hour. One more participant was run, but she showed no standard attentional blink (performance consistently at 100%, except for one lag at 95%) and was therefore excluded. All other participants showed an attentional blink effect (maximum T2 performance − minimum T2 performance) of at least 10%.
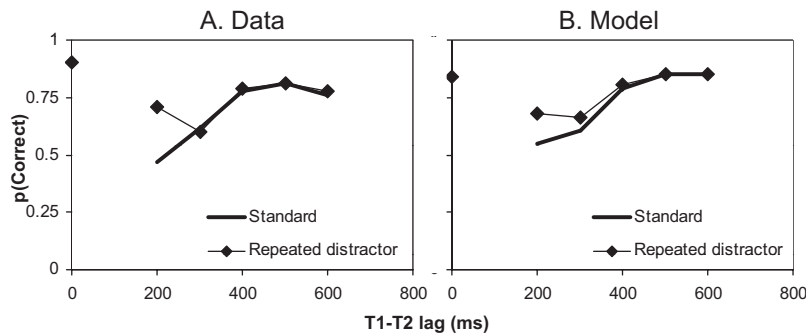
Figure 11. (A) T2 identification performance as a function of T1-T2 lag for two conditions: One in which the first post-T1 distractor has the same identity as the distractor immediately preceding T1 ("repeated") and one in which it has a different identity ("nonrepeated"; equivalent to a standard blink paradigm). Data adapted from "On the Fate of Distractor Stimuli in Rapid Serial Visual Presentation," P. E. Dux, V. Coltheart, & I. Harris, 2006, *Cognition, 99,* 355–382. Copyright 2006 by Elsevier. (B) Model results.
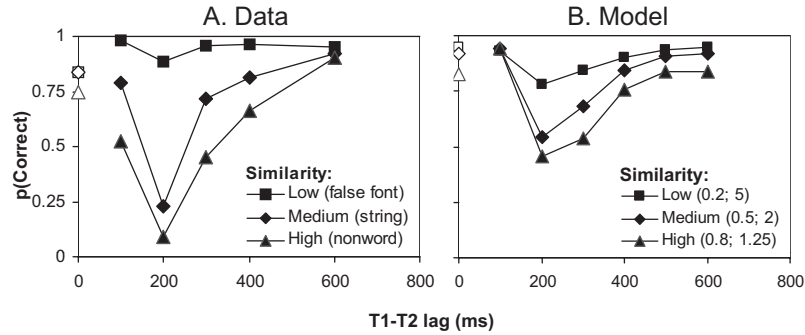
Figure 12. (A) T2 identification performance as a function of T1-T2 lag and as a function of target-distractor similarity, as adapted from "Sources of the Attentional Blink During Rapid Serial Visual Presentation: Perceptual Interference and Retrieval Competition," W. S. Maki, T. Couture, K. Frigen, & D. Lien, 1997, *Journal of Experimental Psychology: Human Perception and Performance, 23,* 1393–1411. Copyright 1997 by the American Psychological Association. (B) Model results, using different values of d_{catdif} and $d_{catsame}$, respectively (which, in our model, stands for the distinctiveness of items between and within categories). The open symbols on the *y* axis signify T1 accuracy.

Stimulus, design, and procedure. Stimulus generation and response recording were done using E-Prime (Psychology Software Tools, Pittsburgh, PA). Backgrounds were gray (40 cd/m²). After a 1000-ms blank period, a $0.5 \times 0.5°$ black fixation cross was presented for 1000 ms in the center of the display and was subsequently replaced by a rapid serial presentation of 24 characters, most of which were black digits, presented in Cou-

rier New (approximately $0.8 \times 0.8°$ in size). Digits were randomly drawn from the set 2–9, with the restriction that no two consecutive digits could be the same. Each item was presented for 67 ms, followed by a 25-ms blank (SOA = 92 ms). At position 8–10 in the stream, a first red letter (T1) replaced one of the black digits. At various lags (1–12), a second red letter (T2) replaced another black digit. The partic-
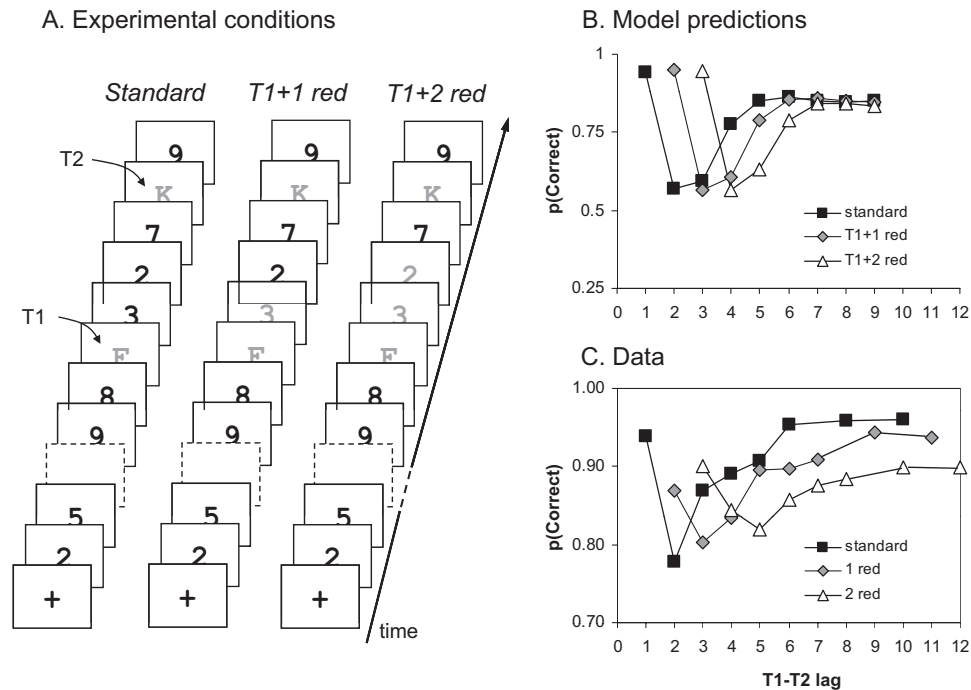
Figure 13. (A) Examples of the three conditions in the new experiment. In the *Standard* condition, targets were red (here drawn in gray) letters among black digit distractors. In the *T1 + 1 red* condition, the distractor immediately following T1 was also red. In the *T1 + 2 red* condition, the two distractors following T1 were red. (B) The boost and bounce model's predictions: The attentional blink is postponed. (C) Average accuracy data from 18 participants for T2 as a function of T1-T2 lag.

ipant's task was to report the red letters at the end of the trial, unspeeded. The letters I, O, Q, and S were excluded. The important manipulation involved the nature of the digit distractors between T1 and T2. In the *standard* condition, these were all black. In the *T1 + 1 red* condition, the digit immediately following T1 was red. In the *T1 + 2 red* condition, the two distractors following T1 were both red. In the standard condition, T2 would follow at lag 1, 2, 3, 4, 5, 6, 8, or 10. In the *T1 + 1 red* condition, T2 could never appear at lag 1 (as this lag was occupied by the red distractor), and thus, lags 2, 3, 4, 5, 6, 7, 9, and 11 were used instead. Similarly, in the *T1 + 2 red* condition, lags were 3, 4, 5, 6, 7, 8, 10, and 12. All conditions were randomly mixed in six blocks of 96 trials each. The first block was regarded as practice. In all, this resulted in 20 trials per cell. The experiment lasted a little more than 1 hr.

Results and Discussion

Mean percentages correct for T1 and T2 (contingent upon T1 correct) were calculated, with order reversals counted as correct. Overall, T1 performance was high at 94%, with little difference across conditions or lag, with the exception that T1 performance dropped to 80% when it was immediately followed by T2 (as was the case only for lag 1 in the standard condition). For other lags and conditions, there were no significant effects, although there was a trend toward deteriorated performance when T1 was followed by red distractors, $F(1, 17) = 3.45$, $MSE = 0.002$, $p = .081$ (all other ps > 0.10). This is to be expected under the assumption that similar distractors act as somewhat stronger masks for T1.

Figure 13c shows T2 performance contingent upon T1 correct. In line with the model's predictions, the pattern of results suggests a clear postponement of the attentional blink with increasing number of post-T1 distractors. In the *standard* condition, sparing occurred for lag 1, whereas the trough of the blink fell at lag 2. In the *T1 + 1 red* condition, relative sparing occurred at lag 2, whereas the trough occurred at lag 3. In the *T1 + 2 red* condition, sparing occurred at lag 3, whereas the trough occurred at lag 5. Because not all lags were the same for all distractor conditions, a single analysis of variance involving the full experimental design was not possible. However, the pattern of results in Figure 13c was statistically confirmed when taking only the common lags into account. An analysis of variance with distractor condition (*standard*, *T1 + 1 red*, *T1 + 2 red*) and the four lags they had in common (lags 3, 4, 5, 6) as factors revealed a significant Distractor Condition × Lag interaction, $F(6, 102) = 5.16$, $p < .001$. The same was true when comparing only the standard to the *T1 + 1 red* condition on all common lags (2, 3, 4, 5, 6), $F(4, 68) = 8.52$, $p < .001$, and when comparing the *T1 + 2 red* with the *T1 + 1 red* condition on all common lags (3, 4, 5, 6, 7), $F(4, 68) = 6.44$, $p < .001$. Finally, the interaction also held when comparing the *T1 + 2 red* condition with the standard condition on common lags (3, 4, 5, 6, 8, 10), $F(5, 85) = 3.82$, $p < .01$. In this latter comparison, there was also a main effect of distractor condition, $F(1, 17) = 21.0$, $p < .001$, indicating that performance was overall worse when T1 was followed by two red distractors.

The data largely confirm the model's predictions. When the post-T1 distractors are at least partly relevant to the gating system, the inhibition is delayed, leading to a postponed attentional blink. As we stated earlier, to qualitatively capture as many data patterns

as possible with a limited set of parameters, we kept parameter values constant for all simulations. With these parameter values, there are some differences between the model's performance and human performance. First, the model shows an overall deeper blink than did human observers.[3] Second, the model shows relatively stronger sparing at the first lag, following the last red distractor. Third, the data show overall subdued (i.e., flattened) attentional blink patterns for the *T1 + 1* and *T1 + 2 red* conditions, whereas this is not the case for the model. All these discrepancies are more due to quantitative than qualitative differences, to the extent that the model assigns relatively more weight to color than humans might, resulting in stronger color-based modulation.

In any case, limited-capacity theories have difficulty explaining the postponement of the attentional blink. Interference theory (Shapiro & Raymond, 1994) predicts stronger interference and, hence, a *deeper* blink from more similar distractors—in this case, the red ones. Two-stage theory sees T1 as the cause of the blink and therefore fails to explain its postponement. If anything, it predicts that red distractors are stronger T1 masks (there was indeed a trend in our data) and, hence, should result in more resources being allocated to T1 at the expense of T2. As mentioned, TLC also predicts a data pattern time-locked to T1, whereas the data suggest that the entire blink curve is shifted backward in time. Finally, it is worth looking at order reversals. To explain the relative sparing at lags 2 and 3 in the red distractor conditions, Bowman and Wyble's (2007) simultaneous type serial token (ST^2) model (as discussed later) might predict that the spared item and T1 are bound to the same token, leading to potential order reversals. However, although our data show substantial order reversals when T2 is spared at lag 1 (23%, in the *standard* condition), there was no sign of increased order reversals in the lag 2 and lag 3 sparing conditions (5% in the *T1 + 1 red* and 3% in the *T1 + 2 red* conditions, relative to 4% for black distractors at the same lags). Thus, sparing does not necessarily go together with order reversals.

General Discussion

We have presented a relatively straightforward theory that explains a considerable number of findings, some of which appeared contradictory at the start. It proposes that target stimuli lead to transient attentional enhancement, whereas distractor stimuli lead to transient attentional suppression. This suppression is particularly strong when distractors are first enhanced, as occurs during the attentional blink.

All of the individual components of our theory have been proposed before in the literature on visual attention. What is new about the theory, however, is the combination of these components and how they interact, allowing for an integrated account of how selection takes place in time. The important novel theoretical step is that selection from rapid visual streams can be accounted for by the interplay of two identical feedback functions, both rapidly rising and more slowly decaying but opposite in sign. As a consequence, and as

---

[3] We fully replicated the experiment with a more difficult set of targets and distractors. Overall, the blink was much deeper, and sparing was reduced; but again, the attentional blink was clearly postponed by red distractors.

a crucial departure from existing theories, there is no central role for limited-capacity resources in our theory. According to Boost and Bounce theory, the biggest piece of evidence for such limitations—the attentional blink—does actually not reflect an extended failure of attention but, rather, its intact, responsive, and rapid operation, modulating the input within about 100 ms. In fact, the attentional blink is caused by too strong an attentional response, rather than by too weak an attentional response. The same strong response accounts for transient attention effects in the cueing studies. Thus, the theory integrates classic transient attention findings with the classic phenomenon of the attentional blink with a relatively limited set of assumptions. It explains numerous associated findings, such as lag-1 sparing, sparing at later lags; the time-based, rather than item-based, nature of the phenomena; apparent trade-offs between T1 and T2; the importance of the distractors; variations in the processing of distractors; similarity effects; and rapid reversals of the blink. By regarding the attentional blink as the direct counterpart of the transient attentional episode, the theory does away with what seemed to be ad hoc explanations of some of these findings. Finally, the theory's computational implementation successfully captures all these effects with a minimal set of parameters with fixed values and generates new predictions. In short, the theory can explain more data with fewer assumptions than most other theories.

There are now a number of theories and computational models of temporal attention, some of which, at the surface, bear similarity to ours. Most of these are theories of the attentional blink. Here we discuss a few of them in detail, whereas others have already been extensively discussed and dismissed by Bowman and Wyble (2007).

Relation to Other Models and Theories: Temporary Suppression Models

As mentioned, boost and bounce theory comes closest to the temporary suppression account originally proposed by Raymond et al (1992). They proposed that T1 initiates an attentional episode that includes the post-T1 distractor. This post-T1 distractor then initiates a period of suppression, during which subsequent items are missed. However, the accounts also differ in important respects:

1. As mentioned earlier, Raymond et al. (1992) hypothesized that the inhibition operates at an early visual stage of processing. To account for the high level representations of blinked items, boost and bounce theory proposes that the inhibition is initiated late, at the level of working memory entrance. However, inhibition does trickle down the representational hierarchy to lower-level representations, affecting the processing of later items.

2. Whereas Raymond et al.'s (1992) account assumed the inhibition to be ballistic and nonadaptive to changes in the stream, in boost and bounce theory, the gating is dynamic and flexible. This allows for the attentional blink to be postponed (see the new experiment here) or rapidly reversed (Di Lollo et al., 2005; Olivers et al., 2007; Kawahara, Kumada, & Di Lollo, 2006; Nieuwenstein et al., 2005).

3. Raymond et al. (1992) proposed that the inhibition is necessary to protect T1 from conjunction errors. As such, it predicts that the blink is time-locked to the T1-versus-first-distractor conflict. In contrast, within boost and bounce theory, T1 does not need to be protected. The theory does not assume that the post-T1 item must be inhibited for T1 to be reported correctly. The inhibition is not triggered by the potential for conjunction errors, conflict, or competition between T1 and the post-T1 item (or between T1 and T2 for that matter) and, thus, is not time-locked to T1. Instead, the inhibition is triggered by a mismatch between a strong incoming signal and the attentional set for the target (see also Olivers & Watson, 2006).

The idea that T1 requires some form of protection has recently returned in a number of verbal and computational models. For example, Hommel et al. (2006) proposed that the post-T1 inhibition is a direct consequence of the attentional network being able to process only one object at a time: The network "silences itself" in an attempt to preserve T1. The idea of a (potential for) conflict between T1 and the post-T1 distractor also returns in Battye's (2006) account of the blink. As we propose here, his model assumes that T1 triggers a temporary enhancement of the post-T1 distractors, causing maximum conflict in the output layers of the network. This conflict is registered and results in stronger lateral inhibition between the output units. T2, when it arrives, suffers from this lateral inhibition, resulting in an attentional blink. As the conflict is being resolved, the blink subsides again. As with Raymond et al.'s (1992) original proposal, one problem with these models is that the inhibition is ballistic and irreversible, contrary to what the data suggest (Olivers et al., 2007; the current experiment). What is more, the idea that higher structures cannot deal with more than one item at a time is highly reminiscent of resource-depletion theories and, thus, suffers from the arguments against these theories we made earlier.

ST^2

A number of other computational models of the attentional blink feature an enhanced response following T1, and many also feature inhibition (Bowman & Wyble, 2007; Chartier et al., 2004; Dehaene et al., 2003; Fragopanagos et al., 2005; Nieuwenhuis et al., 2005). However, in all these models, the enhanced T1 representation either directly inhibits the T2 representation (as a type of lateral inhibition), or it inhibits a type of resource (e.g., working memory capacity or attentional enhancement) that is then not available to T2. These models are therefore actually straightforward limited-capacity models, suffering from the same arguments we made earlier.

Probably the most sophisticated of these models is the ST^2 model of Bowman and Wyble (2007). The ST^2 model is a neural network combining Chun and Potter's (1995) two-stage theory with Kanwisher's (1987) *type/token* distinction. Types are general, context-free representations of objects (such as letters or digits), which may be activated in parallel, without capacity limitations. In contrast, tokens are specific, episodic instances of an object (e.g., the digit "6" presented about midway in the stream). For conscious report, the type representations need to be bound to tokens. This binding occurs in a binding pool. It is important to note that, according to Bowman and Wyble (2007), this tokenization takes time and occurs serially: In principle, only

one token at a time can be bound to a type. Furthermore, the tokenization requires a *blaster*. The blaster is triggered by, and causes attentional enhancement of, T1, so that it can be bound to a token. The blaster also spills over to the next item, so that occasionally, both items may be bound to the same token. This explains lag-1 sparing, and, because temporal information is lost within a token, it also explains the order reversals often found to accompany lag-1 sparing. The attentional blink is explained by assuming that the system seeks to prevent false bindings of post-T1 items to the same token. Therefore, the tokenization process must be temporarily halted. This is done by inhibiting the blaster for a period of several hundreds of milliseconds, until the tokenization of T1 is completed. With a switched-off blaster, subsequent items no longer receive the attentional enhancement necessary for tokenization (and, thus, for conscious report), resulting in an attentional blink. With this architecture, ST^2 can explain an impressive range of data (Bowman & Wyble, 2007).

There are obvious similarities between ST^2 and boost and bounce theory, most notably the presence of both excitatory and inhibitory components. However, there are also fundamental differences:

1. The ST^2 model is, in essence, a hybrid limited-capacity/T1 protection model similar to the ones we discussed above: Tokenization capacity is limited, and to protect the system from false T1 token bindings, subsequent processing is inhibited. Boost and bounce theory argues that there is no need to assume a role for such limitations or protection mechanisms in explaining the attentional blink. A flexible but slightly delayed gating mechanism is sufficient.

2. Within the ST^2 model, the inhibition is triggered by T1 and is ballistic. It therefore cannot easily explain the spread sparing. The ST^2 model may be stretched so that the tokenization process can include more than two successive targets. To account for the Olivers et al. data (2007), Bowman and Wyble (2007) suggested that the blaster may rapidly refire when incoming information remains relevant. As a result, a continuous sequence of targets can be encoded together into working memory (Bowman, Wyble, Chennu, & Craston, 2008; Wyble, Bowman, & Nieuwenstein, in press). This way, the theory might also account for the postponement of the blink, as found in the present experiment. However, it would have to assume that the post-T1 distractors are also tokenized (as they did not inhibit the blaster).

3. It appears to us that even an ongoing blaster within ST^2 could not easily account for rapid reversals of the attentional blink. Here a full blink for a second target was first induced (suggesting that the inhibition must therefore have been triggered), but then a third target escaped the attentional blink when immediately following the second. This goes against the irreversible inhibition of a blaster. Now not only continuous firing but also rapid resumption of the blaster (once it is inhibited) would have to be assumed. The big question regarding Points 2 and 3, then, is why ST^2 would allow for such rapid lifting of the inhibition when it deems the inhibition necessary to pro-

tect T1 in the first place. It appears then that the inhibition protects against something that is not really a threat. In contrast, our boost and bounce model has little trouble with reversing the blink. It assumes that the gate to working memory responds dynamically and adaptively to the incoming sensory information, closing when distractors enter, opening when targets enter.

The Locus Coeruleus Norepinephrine Model

Another model we discuss here is the locus coeruleus norepinephrine (LC-NE) model proposed by Nieuwenhuis et al. (2005). According to this neurocomputational model, the attentional blink is the consequence of the dynamics of the LC, which is the brain stem nucleus responsible for the release of NE in the neocortex. The NE release is believed to have an attention-enhancing effect and, according to Nieuwenhuis et al., is a necessary condition for target awareness. It is interesting to note that LC activity peaks around 100 ms after target onset and, because of auto-inhibition, is then followed by a refractory period of between 200 and 400 ms, during which it cannot fire again (at least in monkeys—little is known about the exact LC dynamics in humans; Aston-Jones, Rajkowski, & Cohen, 2000). According to Nieuwenhuis et al., lag-1 sparing corresponds directly to the peak, whereas the attentional blink corresponds directly to the refractory period. The LC-NE model differs from the boost and bounce model in that the former is, in essence, a limited-capacity resource model. As in the classic attentional blink models, T1 uses up a vital resource that it requires for detection (namely, LC firing and NE release). This resource is then temporarily unavailable to subsequent targets (because of the refractory period). The problem with this is that the LC-NE model cannot explain why the attentional blink is so dependent on the distractors intervening T1 and T2. This is because the LC is triggered by T1, and thus, a refractory period follows, regardless of the presence or nature of the subsequent items. In a similar vein, because it operates in a ballistic fashion, the LC-NE model cannot deal with the sparing of multiple targets and the attentional blink reversal effects mentioned earlier (Di Lollo et al., 2005; Olivers et al., 2007). Nor can it explain the postponement of the blink in the currently presented experiment. Nevertheless, the LC-NE system remains a strong contender for contributing to the transient attentional enhancement effects that are at the core of our model. It is possible that the NE release further modulates the gating of relevant and irrelevant sensory input, thus aggravating the attentional blink (see also Olivers, 2007, for the idea that the LC might be an engine behind transient attentional enhancement).

Gating Theory

The idea that input is filtered, or gated, by working memory has been proposed by many but computationally implemented by few. One of the most prominent mathematical descriptions of how this takes place as a function of time is the gating theory developed by Sperling and colleagues (Reeves & Sperling, 1986; Shih & Sperling, 2002; Sperling & Reeves, 1980; Sperling & Weichselgartner, 1995). According to this theory, the occurrence of a relevant event (e.g., a cue to switch streams or tasks) induces a transition from one discreet attentional state to another. The time course of this transition is described by a temporal transition function, which, on

the basis of Sperling and colleagues' RSVP work, was hypothesized to be gamma-shaped (see Figure 2c). The idea of gating, as well as the rapidly rising and slowly decaying signal (with the shape of a gamma distribution), are at the heart of our model. However, this is largely where the similarities end:

1. Boost and bounce model employs an excitatory as well as an inhibitory gating function. Gating theory mentions only attentional enhancement.

2. Gating theory was designed to describe the time course that accompanies *shifts* or *transitions* in attention, when either a spatial or a task switch is required. As such, it cannot account for the attentional blink, which also occurs when no such switches are necessary. In contrast, the core of boost and bounce theory is designed to account for pure temporal attention phenomena, focusing on situations in which all items appear in the same location and belong to the same task, with the possibility to plug in spatial and task-switch modules (as we have provisionally done in the current implementation).

Recently, to account for the attentional blink, Shih (2008) extended gating theory with a working memory consolidation process that refuses to accept further input while it is busy dealing with a target. Thus, this model, too, is a straightforward limited-capacity resource model, facing the same problems with explaining spread sparing and blink reversals that we saw before. For example, to explain Nieuwenstein and Potter's (2006) whole report benefits, Shih (p. 5) merely assumed a "specific processing strategy," without further explanation.

No Capacity Limitations

The most important difference with existing theories is that boost and bounce theory does away with limited-capacity resources and bottlenecks as core explanations of some classic temporal attention phenomena. Instead, our theory sees the time course of attention as the direct result of the close to (but not quite) optimal operation of run-of-the-mill selection mechanisms that would serve the system just fine under natural circumstances—selection mechanisms that also operate in numerous other attention tasks, such as visual search, partial report, or Stroop-like tasks. Relevant information leads to enhancement, whereas irrelevant information leads to inhibition, all within roughly 100 ms. It is simply that in the artificial case of the RSVP paradigm, these enhancement and inhibition mechanisms are a fraction too late, thus exerting their effects on the next object. Again, we do not deny the existence of capacity limitations. In fact, they play some role in our own theory, when we assume that working memory capacity is limited to about four items (or a little more, when allowing for mechanisms like chunking). This limitation predominantly plays up when more than the maximum number of items enter working memory (e.g., Nieuwenstein & Potter, 2006; Reeves & Sperling, 1986), working memory is already filled by some other task (e.g., Akyürek & Hommel, 2005; Nieuwenstein, Johnson, Kanai, & Martens, 2007), or items in working memory need to be reported in a certain order (Reeves & Sperling, 1986).

Psychological Refractory Period

Capacity limitations or bottlenecks also play an important role in theories of the psychological refractory period (PRP; Pashler, 1984; Welford, 1952). The PRP is the temporary slowing of responses to a second target when that target is presented shortly after a first target. Although the PRP has been linked to the attentional blink (e.g., Jolicoeur, 1998), we believe there are important methodological differences that would allow for two independent effects. For one, within the PRP paradigm, observers usually need to respond immediately and as quickly as possible to both targets. Nevertheless, observers are typically fully aware of both targets. This leaves open the possibility that the PRP is the consequence of competition on the response level (reflecting a true bottleneck), whereas in our view, the attentional blink reflects the active gating of working memory access, even before this bottleneck can come into play (as the prerequisites for generating a response are not even there). Second, PRP experiments typically involve a clear task switch and, often, a sensory modality switch. Such switches may be regarded as a contamination of the attentional blink (as the blink also occurs perfectly without task-switching components; Chun & Potter, 2001; Enns, Visser, Kawahara, & Di Lollo, 2001; Potter, Chun, Banks, & Muckenhoupt, 1998). Third, the PRP paradigm is usually devoid of distractors, whereas it has been shown that distractors play an important role in the attentional blink paradigm. In all, then, the PRP appears to reflect processes that occur after selection, whereas in the present article, we are interested in the time course of selection itself (and we argue that the attentional blink reflects just that). Marois and Ivanoff (2005) have mentioned a few more grounds on which the PRP and the attentional blink might be dissociated. Naturally, however, our model could be extended with components important for explaining PRP phenomena.

But What About Previously Reported Resource-Depletion Effects in the Attentional Blink Paradigm?

Limited-capacity theories predict that T1 and T2 act like two communicating vessels: The more resources T1 demands, the fewer are left for T2. Surely there is evidence for that? Well, in our view, the evidence is actually rather thin. A number of studies have indeed found that more difficult T1s lead to deeper blinks for T2 (Grandison et al., 1997; Jolicoeur, 1998; Ouimet & Jolicoeur, 2007; Seiffert & Di Lollo, 1997; Shore, McLaughlin, & Klein, 2001; Visser, 2007). However, as also pointed out by McLaughlin, Shore, and Klein (2001), many of the experimental procedures in these studies involved a task switch from T1 to T2 (and some even involved a location switch). Switch costs may be aggravated by a difficult T1 task, whereas the attentional blink itself is not. Some experiments required participants to respond immediately to T1, which, together with the task switch, turns the procedure into a PRP paradigm, with possible bottlenecks at the response-selection level (see the previous section). Moreover, others have changed the post-T1 mask as a way to make T1 more difficult to identify. However, also in boost and bounce theory, a stronger mask is predicted to result in a stronger blink, simply because it induces stronger inhibition. As argued earlier, differential masking can also partly explain the T1/T2 tradeoff found by Potter et al. (2002).

Other studies failed to find a tradeoff altogether (McLaughlin et al., 2001; Shapiro et al., 1994; Ward et al., 1997) or reported only

additive effects (i.e., T1 load did not interact with lag or presentation time, suggesting that it affects stages independent of what causes the attentional blink; Akyürek & Hommel, 2005; Jolicoeur & Dell'Acqua, 2000; Nieuwenstein et al., 2007; Olivers et al., 2007). Again, others have reported an opposite tradeoff (deeper blinks after easier T1s; Chua, 2005).

Conversely, T2 performance can be improved without impeding T1 performance. Participants in our lab occasionally reported that they felt they did better when they did not focus so strongly on the stimulus stream, an observation also made earlier by Kahneman (1973). Olivers and Nieuwenhuis (2005, 2006; see also Arend, Johnston, & Shapiro, 2006, and Ho, Mason, & Spence, 2006) put this observation to the test by comparing performance in a standard attentional blink condition, in which participants were instructed to concentrate on the task, with performance in conditions in which observers were somewhat distracted from the central task. For example, we asked them to actively think about their holiday plans, to listen to a repetitive tune and detect an occasional yell in it, to perform an additional memory task, or to simply concentrate a little less. All these manipulations had the same effect: Detection of the second target improved under conditions intended to invoke a more distributed attentional state, suggesting that taking away attentional resources may actually be beneficial, rather than detrimental (but see T. Spalek & V. Di Lollo, personal communication, summer 2005, for a failure to replicate). Similarly, using multiple RSVP streams, Kristjánsson and Nakayama (2002) found the attentional blink to be more reduced the further away T2 was from the stream containing T1 and, thus, the fewer resources were presumably allocated to the T2 stream. Furthermore, Slagter et al. (2007) found that extensive meditation practice leads to a reduced attentional blink. It is important to note that in none of these studies did the improvement in T2 performance go at the expense of T1 performance.

In a similar vein, a study by Ferlazzo, Lucido, Di Nocera, Fagioli, and Sdoia (2007) also suggests that processing resources are not necessarily limited to one target, at the expense of the other target. They found a standard attentional blink for the second of two targets when participants were instructed to report "each of the two targets." However, using identical RSVPs, they found a reduced blink when the instruction was to report "the pair of targets." The blink was even virtually absent when the instructions were to report "the sum of the targets" (when they were digits) or "the syllable formed by the targets" (when they were letters). Ferlazzo et al. argued that normally, participants treat T1 and T2 report as separate goals, between which they then need to switch. It is this goal switch that causes the attentional blink. When participants are induced to adopt a more holistic, distributed task set, the attentional blink is reduced or absent. Finally, the fact that the attentional blink disappears when the RSVP stream can be treated as a single, gradually changing object (Raymond, 2003; Kellie & Shapiro, 2004) also points toward the conclusion that what matters is the selection mechanism itself (i.e., whether or not individual items need to be selected from the stream) and not so much what happens to T1 after selection.

Whereas limited-capacity theories have difficulties explaining these results, our theory readily explains why there is no unequivocal evidence for resource depletion or structural bottlenecks in the attentional blink. This is because resource depletion plays no central role in our theory. Moreover, a more holistic or distributed approach toward the stream would mean a weaker inhibitory response to the distractors and, thus, a weaker blink, simply because distractors are not treated as such (or not as strongly).

Neurophysiology

The attentional mechanisms we have proposed here are consistent with what is known about the nervous system. The idea that selection takes place through a rapid spread of recurrent processing down the representational hierarchy has been proposed before (Di Lollo et al., 2000) and has been directly linked to neurophysiological processes (Dehaene et al., 1998; Grossberg, 1995; Lamme & Roelfsema, 2000). There is also ample evidence that such recurrent processes might be biased in advance of stimulus presentation, effectively allowing for an attentional filter or template to be set up (e.g., Chelazzi, Duncan, Miller, & Desimone, 1998; Desimone & Duncan, 1995; Motter, 1994). A likely source for such biasing or gating signals is prefrontal cortex, an area linked to working memory function (E. K. Miller, Erickson, & Desimone, 1996; E. K. Miller & Cohen, 2001). The prefrontal cortex has long-range feedback connections to extrastriate visual areas as well as the inferotemporal cortex, making it the prime candidate for hosting the gate neurons that we assume modulate the rise and fall of attention in response to incoming stimuli. A recent study by McNab and Klinberg (2008) is especially interesting here. In their functional magnetic resonance imaging study, they found the posterior part of the middle frontal gyrus to become active when observers were instructed to expect distractors during a subsequent working memory task (without the stimuli being present yet), indicating that this area represents the task settings necessary to reject distractors. Activity in another structure, the basal ganglia, was also increased prior to distractor rejection and correlated negatively with the unnecessary storage of items in working memory, as indicated by parietal activity (Vogel, McCollough, & Machizawa, 2005). This led Awh and Vogel (2008) to call the prefrontal cortex/basal ganglia network "the bouncer in the brain." It is interesting to note that the basal ganglia, in turn, have been proposed to provide the transient boost of activity necessary to update working memory (Hazy et al., 2006), making it another serious candidate for being an engine behind transient attention in our model.

Further evidence comes from electroencephalography (EEG) studies. The frontal cortex is also the source of the event-related potential known as the frontal selection positivity (FSP; Potts, 2004). This component is known to respond selectively to task-relevant stimuli and may correlate with the gating mechanisms governing access to working memory. Two EEG studies using attentional blink tasks show an encouraging pattern of findings in this respect. At frontal electrode sites, Martens, Munneke, et al. (2006) found a marked positivity roughly 250 ms after T1 onset, relative to a condition in which there was no T1. It is interesting to note that, after approximately 100–150 ms, the FSP was followed by a marked and prolonged negativity in the EEG signal, relative to the no-target baseline. This negativity reached its peak shortly after the FSP, gradually returning to baseline over a period of 300–500 ms. Niedeggen, Hesselman, Sharaie, Milders, and Blakemore (2004) reported a highly similar pattern and indeed linked it to a "postperceptual frontal gating mechanism that controls the

access of visual stimuli to higher order evaluation" (p. 584). We believe that this post-FSP negativity may reflect a direct neurophysiological correlate of the attentional blink, in terms of locus as well as time course.[4]

Another EEG component that appears diagnostic of the attentional blink is the P3/P300, with the M300 as its magnetoencephalography counterpart. The P3 is a broad, positive event-related brain potential probably encompassing multiple components but with a modal peak latency of about 300–350 ms after stimulus onset. It has often been associated with the updating of working memory (Donchin & Coles, 1988; Vogel & Luck, 2002). With regard to the attentional blink, there is a positive correlation between T2 accuracy and the amplitude and/or latency of the P3 elicited by T2 (Kessler et al., 2005; Rolke et al., 2001; Sergent et al., 2005; Vogel et al., 1998). At the same time, T2 accuracy has also been found to correlate negatively with the amplitude of the P3 elicited by T1 (Martens, Johnson, Elmallah, & London, 2006; McArthur, Budd, & Michie, 1999; Shapiro, Schmitz, Martens, Hommel, & Schnitzler, 2006). These findings appear to lend support to the idea of a trade-off in resources between T1 and T2 on the level of working memory. However, one needs to keep in mind that these relationships are purely correlational, not causal. This means that a third factor may explain the negative correlation between the P3 to T1 and the P3 to T2. For example, imagine the following scenario: On a particular trial, T1 is detected and elicits a particularly strong attentional enhancement, which is then measured as a strong P3. The strong enhancement also carries over to the post-T1 distractor, with an equally strong inhibitory response as a result. Because of this strong suppression, T2 is less likely to enter working memory and is, hence, less likely to trigger a P3. This explains why strong blinks correlate with strong P3s to T1 and weak P3s to T2. Thus, in this scenario, the negative correlation between T1 and T2 does not reflect a direct tradeoff in resources between the two targets but, rather, a common underlying third factor, namely a strong attentional boost. At present, excitatory feedback in our model is of fixed strength and length, which is sufficient to account for averaged data. Naturally, it could be made variable to account for differences between individual trials.

Conclusion

Boost and bounce theory comprehensively and coherently explains the time course of attention and associated phenomena in number of different paradigms involving rapid serial processing. Moreover, the theory has been made explicit in a working computational model based on a minimum of assumptions and makes testable new predictions. Finally, it appears to connect well with current knowledge of the neurophysiology of attention. There is no doubt that, as a theory of attention, the theory is far from complete. Notably, the interactions with spatial and feature processing need to be made explicit, as do the higher-end mechanisms involved in the setting up and switching of tasks and the selection and consolidation of response representations. What the present work aims to demonstrate is that the dynamics of attention can largely be captured by a straightforward gating system, thus fulfilling attention's classic role as the gatekeeper between sensory and response processes.

---

[4] The linking of inhibitory processing to a negative component is accidental, as the polarity of event-related potential components is a function of brain anatomy, not physiology.

References

Akyürek, E. G., & Hommel, B. (2005). Short-term memory and the attentional blink: Capacity versus content. *Memory & Cognition, 33,* 654–663.

Anderson, A. K. (2005). Affective influences on the attentional dynamics supporting awareness. *Journal of Experimental Psychology: General, 134,* 258–281.

Arend, I., Johnston, S., & Shapiro, K. (2006). Task-irrelevant visual motion and flicker attenuate the attentional blink. *Psychonomic Bulletin & Review, 13,* 600–607.

Aston-Jones, G., Rajkowski, J., & Cohen, J. D. (2000). Locus coeruleus and regulation of behavioral flexibility and attention. *Progress in Brain Research, 126,* 549–554.

Awh, E., & Vogel, E. K. (2008). The bouncer in the brain. *Nature Neuroscience, 11,* 5–6.

Baars, B. J. (1989). *A cognitive theory of consciousness.* Cambridge, United Kingdom: Cambridge University Press.

Bachmann, T., & Hommuk, K. (2005). How backward masking becomes attentional blink. *Psychological Science, 16,* 740–742.

Bachmann, T., & Oja, A. (2003). Flash-lag without change in feature space is alive and well at late intervals after stream onset. *Perception, S32,* 126–127.

Baddeley, A. D., & Hitch, G. (1974). Working memory. In G. Bower (Ed.), *Recent advances in learning and motivation* (Vol. 8, pp. 47–90). New York: Academic Press.

Bargh, J. (1992). The ecology of automaticity: Toward establishing the conditions needed to produce automatic processing effects. *American Journal of Psychology, 105*(2), 181–199.

Battye, T. G. G. (2006). *Connectionist modelling of attention and anxiety.* Unpublished PhD thesis, The Medical Research Council's Cognition and Brain Sciences Unit, Cambridge University, United Kingdom.

Botella, J., Barriopedro, M., & Suero, M. (2001). A model of the formation of illusory conjunctions in the time domain. *Journal of Experimental Psychology: Human Perception and Performance, 27,* 1452–1467.

Bowman, H., & Wyble, B. P. (2007). The simultaneous type, serial token model of temporal attention and working memory. *Psychological Review, 114,* 38–70.

Bowman, H., Wyble, B., Chennu, S., & Craston, P. (2008). A reciprocal relationship between bottom-up trace strength and the attentional blink bottleneck: Relating the LC–NE and ST2 Models. *Brain Research, 1202,* 25–42.

Brehaut, J. C., Enns, J. T., & Di Lollo, V. (1999). Visual masking plays two roles in the attentional blink. *Perception & Psychophysics, 61*(7), 1436–1448.

Breitmeyer, B., Ehrenstein, A., Pritchard, K., Hiscock, M., & Crisan, J. (1999). The roles of location specificity and masking mechanisms in the attentional blink. *Perception & Psychophysics, 61*(5), 798–809.

Breitmeyer, B. G., & Ganz, L. (1976). Implications of sustained and transient channels for theories of visual pattern masking, saccadic suppression, and information processing. *Psychological Review, 83,* 1–36.

Broadbent, D. E., & Broadbent, M. H. P. (1987). From detection to identification: Response to multiple targets in rapid serial visual presentation. *Perception & Psychophysics, 42,* 105–113.

Bundesen, C., Habekost, T., & Kyllingsbæk, S. (2005). A neural theory of visual attention: Bridging cognition and neurophysiology. *Psychological Review, 112,* 291–328.

Busey, T. A., & Loftus, G. R. (1994). Sensory and cognitive components of visual informaiton acquisition. *Psychological Review, 101,* 446–469.

Carlson, T. A., Hogendoorn, H., & Verstraten, F. A. J. (2006). The speed of visual attention: What time is it? *Journal of Vision, 6,* 1406–1411.

Chartier, S., Cousineau, D., & Charbonneau, D. (2004). A connectionist model of the attentional blink effect during a rapid serial visual task. *Proceedings of the 6th International Conference on Cognitive Modeling,* 64–69.

Cheal, M. L., & Lyon, D. R. (1991). Central and peripheral precuing of forced-choice discrimination. *Quarterly Journal of Experimental Psychology, 43*A, 859–880.

Chelazzi, L., Duncan, J., Miller, E. K., & Desimone, R. (1998). Responses of neurons in inferior temporal cortex during memory-guided visual search. *Journal of Neurophysiology, 80,* 2918–2940.

Chua, F. K. (2005). The effect of target contrast on the attentional blink. *Perception & Psychophysics, 67,* 770–788.

Chua, F. K., Goh, J., & Hon, N. (2001). Nature of codes extracted during the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance, 27,* 1229–1242.

Chun, M. M. (1997a). Temporal binding errors are redistributed by the attentional blink. *Perception & Psychophysics, 59*(8), 1191–1199.

Chun, M. M. (1997b). Types and tokens in visual processing: A double dissociation between the attentional blink and repetition blindness. *Journal of Experimental Psychology: Human Perception and Performance, 23,* 738–755.

Chun, M. M., & Potter, M. C. (1995). A two-stage model for multiple detection in RSVP. *Journal of Experimental Psychology: Human Perception and Performance, 21,* 109–127.

Chun, M. M., & Potter, M. C. (2001). The attentional blink and task switching within and across modalities. In K. Shapiro (Ed.), *The limits of attention: Temporal constraints in human information processing* (pp. 20–35). London: Oxford University Press.

Connors, B. W., & Gutnick, M. J. (1990). Intrinsic firing patterns of diverse neocortical neurons. *Trends in Neuroscience, 13,* 99–104.

Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences, 24,* 87–114.

Dehaene, S., Kerszberg, M., & Changeux, J.-P. (1998). A neuronal model of a global workspace in effortful cognitive tasks. *Proceedings of the National Academy of Sciences, USA, 95,* 14529–14534.

Dehaene, S., Sergent, C., & Changeux, J.-P. (2003). A neuronal network model linking subjective reports and objective physiological data during conscious perception. *Proceedings of the National Academy of Sciences, USA, 100,* 8520–8525.

Dell'Acqua, R., Pascali, A., Jolicoeur, P., & Sessa, P. (2003). Four-dot masking produces the attentional blink. *Vision Research, 43,* 1907–1913.

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

Di Lollo, V., Enns, J. T., & Rensink, R. A. (2000). Competition for consciousness among visual events: The psychophysics of reentrant visual processes. *Journal of Experimental Psychology: General, 129,* 481–507.

Di Lollo, V., Kawahara, J., Ghorashi, S. M. S., & Enns, J. T. (2005). The attentional blink: Resource depletion or temporary loss of control? *Psychological Research, 69,* 191–200.

Donchin, E., & Coles, M. G. H. (1988). Is the P300 component a manifestation of context updating? *Behavioral and Brain Sciences, 11,* 357–374.

Drew, T., & Shapiro, K. (2006). Representational masking and the attentional blink. *Visual Cognition, 13,* 513–528.

Duncan, J., Ward, R., & Shapiro, K. (1994). Direct measurement of attentional dwell time in human vision. *Nature, 369,* 313–315.

Dux, P. E., & Coltheart, V. (2005). The meaning of the mask matters: Evidence of conceptual interferences in the attentional blink. *Psychological Science, 16,* 775–779.

Dux, P. E., Coltheart, V., & Harris, I. (2006). On the fate of distractor stimuli in rapid serial visual presentation. *Cognition, 99,* 355–382.

Engle, R. W., Conway, A. R. A., Tuholski, S. W., & Schisler, R. J. (1995). A resource account of inhibition. *Psychological Science, 6,* 122–125.

Enns, J. T., Visser, T. A. W., Kawahara, J.-I., & Di Lollo, V. (2001). Visual masking and task switching in the attentional blink. In K. Shapiro (Ed.), *The limits of attention: Temporal constraints in human information processing* (pp. 65–81). London: Oxford University Press.

Eriksen, C. W. (1990). Attentional search of the visual field. In D. Brogan (Ed.), *Visual search* (pp. 3–19). London: Taylor & Francis.

Ferlazzo, F., Lucido, S., Di Nocera, F., Fagioli, S., & Sdoia, S. (2007). Switching between goals mediates the attentional blink effect. *Experimental Psychology, 54,* 89–98.

Folk, C., Remington, R. W., & Johnston, J. C. (1992). Involuntary covert orienting is contingent on attentional control settings. *Journal of Experimental Psychology: Human Perception and Performance, 18,* 1030–1044.

Fragopanagos, N., Kockelkoren, S., & Taylor, J. G. (2005). A neurodynamic model of the attentional blink. *Cognitive Brain Research, 24,* 568–586.

Fuster, J. M. (1997). *The prefrontal cortex: Anatomy, physiology, and neuropsychology of the frontal lobe* (3rd ed.). New York: Lippincott Williams & Wilkins.

Fuster, J. M., Bauer, R. H., & Jervey, J. P. (1982). Cellular discharge in the dorsolateral prefrontal cortex of the monkey in cognitive tasks. *Experimental Neurology, 77,* 679–694.

Gawne, T. J., & Martin, J. M. (2002). Responses of primate visual cortical V4 neurons to simultaneously presented stimuli. *Journal of Neurophysiology, 88,* 1128–1135.

Ghorashi, S. M. S., Zuvic, S. M., Visser, T. A. W., & Di Lollo, V. (2003). Focal distraction: Spatial shifts of attention are not required for contingent capture. *Journal of Experimental Psychology: Human Perception and Performance, 24,* 699–706.

Giesbrecht, B., Bischof, W. F., & Kingstone, A. (2003). Visual masking during the attentional blink: Tests of the object substitution hypothesis. *Journal of Experimental Psychology: Human Perception and Performance, 29,* 238–258.

Giesbrecht, B., & Di Lollo, V. (1998). Beyond the attentional blink: Visual masking by object substitution. *Journal of Experimental Psychology: Human Perception and Performance, 24,* 1454–1466.

Goldman-Rakic, P. S. (1995). Cellular basis of working memory. *Neuron, 14,* 477–485.

Grandison, T. D., Ghirardelli, T. G., & Egeth, H. E. (1997). Beyond similarity: Masking of the target is sufficient to cause the attentional blink. *Perception & Psychophysics, 59*(2), 266–274.

Grossberg, S. (1995). The attentive brain. *American Scientist, 83,* 438–449.

Hasher, L., & Zacks, R. T. (1988). Working memory, comprehension, and aging. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 22, pp. 193–225). San Diego, CA: Academic Press.

Hazy, T. E., Frank, M. J., & O'Reilly, R. (2006). Banishing the homunculus: Making working memory work. *Neuroscience, 139,* 105–118.

Hikosaka, O., Miyauchi, S., & Shimojo, S. (1993). Focal visual attention produces illusory temporal order and motion sensation. *Vision Research, 33,* 1219–1240.

Ho, C., Mason, O., & Spence, C. (2006). An investigation into the temporal dimension of the Mozart effect: Evidence from the attentional blink task. *Acta Psychologica, 125,* 117–128.

Hommel, B., & Akyürek, E. G. (2005). Lag-1 sparing in the attentional blink: Benefits and costs of integrating two events into a single episode. *Quarterly Journal of Experimental Psychology, 58,* 1415–1433.

Hommel, B., Kessler, K., Schmitz, F., Gross, J., Akyürek, E. G., Shapiro, K., et al. (2006). How the brain blinks: Towards a neurocognitive model of the attentional blink. *Psychological Research, 70,* 425–435.

Houghton, D. (1994). A model of inhibitory mechanisms in selective

attention. In D. Dagenbach (Ed.), *Inhibitory processes in attention, memory, and language* (pp. 53–112). San Diego, CA: Academic Press.

Huber, D. E., & O'Reilly, R. C. (2003). Persistence and accommodation in short-term priming and other perceptual paradigms: Temporal segregation through synaptic depression. *Cognitive Science, 27,* 403–430.

Isaak, M. I., Shapiro, K. L., & Martin, J. (1999). The attentional blink reflects retrieval competition among multiple rapid serial visual presentation items: Tests of an interference model. *Journal of Experimental Psychology: Human Perception and Performance, 25,* 1774–1792.

Jolicoeur, P. (1998). Modulation of the attentional blink by on-line response selection: Evidence from speeded and unspeeded task-1 decisions. *Memory & Cognition, 26*(5), 1014–1032.

Jolicoeur, P. (1999). Concurrent response-selection demands modulate the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance, 25,* 1097–1113.

Jolicoeur, P., & Dell'Acqua, R. (1998). The demonstration of short-term consolidation. *Cognitive Psychology, 32,* 138–202.

Jolicoeur, P., & Dell'Acqua, R. (2000). Selective influence of second target exposure duration and Task1 load effects in the attentional blink phenomenon. *Psychonomic Bulletin & Review, 7,* 472–479.

Jolicoeur, P., Tombu, M., Oriet, C., & Stevanovski, B. (2002). From perception to action: Making the connection. In W. Prinz & B. Hommel (Eds.), *Attention and performance: Vol. XIX. Common mechanisms in perception and action* (pp. 558–586). Oxford, United Kingdom: Oxford University Press.

Jonides, J. (1981). Voluntary versus automatic control over the mind's eye's movement. In J. Long & A. Baddeley (Eds.), *Attention and performance IX* (pp. 187–203). Hillsdale, NJ: Erlbaum.

Jonides, J., Smith, E. E., Marshuetz, C., Koeppe, R. A., & Reuter-Lorenz, P. A. (1998). Inhibition in verbal working memory revealed by brain activation. *Proceedings of the National Academy of Sciences, USA, 95,* 8410–8413.

Kahneman, D. (1973). *Attention and effort.* New York: Prentice Hall.

Kanwisher, N. (1987). Repetition blindness: Type recognition without token individuation. *Cognition, 27,* 117–143.

Kawahara, J.-I., & Enns, J. T. (in press). Selection difficulty and inter-item competition are independent factors in rapid visual stream perception. *Journal of Experimental Psychology: Human Perception and Performance.*

Kawahara, J., Enns, J. T., & Di Lollo, V. (2006). The attentional blink is not a unitary phenomenon. *Psychological Research, 70,* 405–413.

Kawahara, J., Kumada, T., & Di Lollo, V. (2006). The attentional blink is governed by a temporary loss of control. *Psychonomic Bulletin & Review, 13,* 886–890.

Kellie, F. J., & Shapiro, K. (2004). Object file continuity predicts attentional blink magnitude. *Perception & Psychophysics, 66,* 692–712.

Kessler, K., Schmitz, F., Gross, J., Hommel, B., Shapiro, K., & Schnitzler, A. (2005). Cortical mechanisms of attention in time: Neural correlates of the lag-1 sparing phenomenon. *European Journal of Neuroscience, 21,* 2563–2574.

Keysers, C., & Perrett, D. I. (2002). Visual masking and RSVP reveal neural competition. *Trends in Cognitive Sciences, 6,* 120–125.

Keysers, C., Xiao, D.-K., Földiák, P., & Perrett, D. I. (2001). The speed of sight. *Journal of Cognitive Neuroscience, 13*(1), 90–101.

Kristjánsson, Á., Mackeben, M., & Nakayama, K. (2001). Rapid, object-based learning in the deployment of transient attention. *Perception, 30,* 1375–1387.

Kristjánsson, Á., & Nakayama, K. (2002). The attentional blink in space and time. *Vision Research, 42,* 2039–2050.

Lamme, V. A. F., & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neuroscience, 23,* 571–579.

Lavie, N. (2000). Selective attention and cognitive control: Dissociating attentional functions through different types of load. In S. Monsell & J. Driver (Eds.), *Control of cognitive processes: Attention and performance XVIII* (pp. 175–194). Cambridge, MA: Bradford, MIT Press.

Lavie, N., Hirst, A., Fockert, J. W. d., & Viding, E. (2004). Load theory of selective attention and cognitive control. *Journal of Experimental Psychology: General, 133,* 339–354.

Lawrence, D. H. (1971). Two studies of visual search for word targets with controlled rates of presentation. *Perception and Psychophysics, 10,* 85–89.

Legge, G. (1978). Sustained and transient mechanisms in human vision: Temporal and spatial properties. *Vision Research, 18,* 341–376.

Loach, D., & Marí-Beffa, P. (2003). Post-target inhibition: A temporal binding mechanism? *Visual Cognition, 10,* 513–526.

Luck, S. J., & Vecera, S. P. (2002). Attention: From tasks to mechanisms. In S. Yantis (Ed.), *Stevens' handbook of experimental psychology: Vol. 1. Sensation and perception* (pp. 235–286). New York: Wiley.

Luck, S. J., Vogel, E. K., & Shapiro, K. L. (1996, October 17). Word meanings can be accessed but not reported during the attentional blink. *Nature, 383,* 616–618.

Mackeben, M., & Nakayama, K. (1993). Express attentional shifts. *Vision Research, 33,* 85–90.

Macknik, S. L., & Livingstone, M. S. (1998). Neuronal correlates of visibility and invisibility in the primate visual system. *Nature Neuroscience, 1,* 144–149.

Maki, W. S., Bussard, G., Lopez, K., & Digby, B. (2003). Sources of interference in the attentional blink: Target-distractor similarity revisited. *Perception & Psychophysics, 65,* 188–201.

Maki, W. S., Couture, T., Frigen, K., & Lien, D. (1997). Sources of the attentional blink during rapid serial visual presentation: Perceptual interference and retrieval competition. *Journal of Experimental Psychology: Human Perception and Performance, 23,* 1393–1411.

Maki, W. S., Frigen, K., & Paulson, K. (1997). Associative priming by targets and distractors during rapid serial visual presentation: Does word meaning survive the attentional blink? *Journal of Experimental Psychology: Human Perception and Performance, 23,* 1014–1034.

Maki, W. S., & Padmanabhan, G. (1994). Transient suppression of processing during rapid serial visual presentation: Acquired distinctiveness of probes modulates the attentional blink. *Psychonomic Bulletin & Review, 1,* 499–504.

Marois, R., & Ivanoff, J. (2005). Capacity limits of information processing in the brain. *Trends in Cognitive Sciences, 9,* 296–305.

Marois, R., Yi, D. J., & Chun, M. M. (2004). The neural fate of consciously perceived and missed events in the attentional blink. *Neuron, 41,* 465–472.

Martens, S., Johnson, A., Elmallah, K., & London, R. (2006). Cuing and stimulus effects on the P3 and the AB. *Acta Psychologica, 123,* 204–218.

Martens, S., Munneke, J., Smid, H., & Johnson, A. (2006). Quick minds don't blink: Electrophysiological correlates of individual differences in attentional selection. *Journal of Cognitive Neuroscience, 18,* 1423–1438.

Martens, S., Wolters, G., & Van Raamsdonk, M. (2002). Blinks of the mind: Memory effects of attentional processes. *Journal of Experimental Psychology: Human Perception and Performance, 28,* 1275–1287.

Maylor, E. A., & Hockey, R. (1985). Inhibitory component of externally controlled covert orienting in visual space. *Journal of Experimental Psychology: Human Perception and Performance, 11,* 777–787.

McArthur, G., Budd, T., & Michie, P. (1999). The attentional blink and P300. *NeuroReport, 10*(17), 3691–3695.

McAuliffe, S. P., & Knowlton, B. J. (2000). Dissociating the effects of featural and conceptual interference on multiple target processing in rapid serial visual presentation. *Perception & Psychophysics, 62,* 187–195.

McLaughlin, E. N., Shore, D. I., & Klein, R. M. (2001). The attentional

blink is immune to masking-induced data limits. *Quarterly Journal of Experimental Psychology, 54*A, 169–196.

McNab, F., & Klingberg, T. (2008). Prefrontal cortex and basal ganglia control access to working memory. *Nature Neuroscience, 11,* 103–107.

Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience, 24,* 167–202.

Miller, E. K., Erickson, C. A., & Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *Journal of Neuroscience, 16,* 5154–5167.

Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review, 63,* 81–97.

Miller, J. (1989). The control of attention by abrupt visual onsets and offsets. *Perception & Psychophysics, 45,* 567–571.

Motter, B. C. (1994). Neural correlates of attentive selection for color or luminance in extrastriate area V4. *Journal of Neuroscience, 14,* 2178–2189.

Müller, H. J., & Rabbitt, P. M. A. (1989). Reflexive and voluntary orienting of visual attention: Time course of activation and resistance to interruption. *Journal of Experimental Psychology: Human Perception and Performance, 15,* 315–330.

Nakayama, K., & Mackeben, M. (1989). Sustained and transient components of focal visual attention. *Vision Research, 29,* 1631–1647.

Niedeggen, M., Hesselman, G., Sharaie, A., Milders, M., & Blakemore, C. (2004). Probing the prerequisites for motion blindness. *Journal of Cognitive Neuroscience, 16,* 584–597.

Nieuwenhuis, S., Gilzenrat, M. S., Holmes, B. D., & Cohen, J. D. (2005). The role of the locus coeruleus in mediating the attentional blink: A neurocomputational theory. *Journal of Experimental Psychology: General, 134,* 291–307.

Nieuwenstein, M. R. (2006). Top-down controlled, delayed selection in the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance, 32,* 973–985.

Nieuwenstein, M. R., Chun, M. M., Lubbe, R. H. J. v. d., & Hooge, I. T. C. (2005). Delayed attentional engagement in the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance, 31,* 1463–1475.

Nieuwenstein, M. R., Johnson, A., Kanai, R., & Martens, S. (2007). Cross-task repetition amnesia: Impaired recall of RSVP targets held in memory for a secondary task. *Acta Psychologica, 125,* 319–331.

Nieuwenstein, M. R., & Potter, M. C. (2006). Temporal limits of selection and memory encoding: A comparison of whole versus partial report in rapid serial visual presentation. *Psychological Science, 17,* 471–475.

Nothdurft, H.-C. (2002). Attention shifts to salient targets. *Vision Research, 42,* 1287–1306.

Olivers, C. N. L. (2007). The time course of attention: It's better than we thought. *Current Directions in Psychological Science, 16,* 11–15.

Olivers, C. N. L., & Nieuwenhuis, S. (2005). The beneficial effect of concurrent task-irrelevant mental activity on temporal attention. *Psychological Science, 16*(4), 265–269.

Olivers, C. N. L., & Nieuwenhuis, S. (2006). The beneficial effects of additional task load, positive affect, and instruction on the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance, 32,* 364–379.

Olivers, C. N. L., Van der Stigchel, S., & Hulleman, J. (2007). Spreading the sparing: Against a limited-capacity account of the attentional blink. *Psychological Research, 71,* 126–139.

Olivers, C. N. L., & Watson, D. G. (2006). Input control processes in rapid serial visual presentation: Target selection and distractor inhibition. *Journal of Experimental Psychology: Human Perception and Performance, 32,* 1083–1092.

Ouimet, C., & Jolicoeur, P. (2007). Beyond Task 1 difficulty: The duration of T1 encoding modulates the attentional blink. *Visual Cognition, 15,* 290–304.

Pashler, H. (1984). Processing stages in overlapping tasks: Evidence for a central bottleneck. *Journal of Experimental Psychology: Human Perception and Performance, 10,* 358–377.

Popple, A. V., & Levi, D. M. (2007). Attentional blinks as errors in temporal binding. *Vision Research, 47,* 2973–2981.

Posner, M. I. (1980). Orienting of attention: The 7th Sir Frederic Bartlett lecture. *Quarterly Journal of Experimental Psychology, 32,* 3–25.

Posner, M. I., & Cohen, Y. (1984). Components of visual orienting. In H. Bouma & D. G. Bouwhuis (Eds.), *Attention and performance X: Control of language processes* (pp. 531–556). Hillsdale, NJ: Erlbaum.

Postle, B. R. (2006). Working memory as an emergent property of the mind and brain. *Neuroscience, 139,* 23–38.

Potter, M. C., Chun, M. M., Banks, B. S., & Muckenhoupt, M. (1998). Two attentional deficits in serial target search: The visual attentional blink and an amodal task-switch deficit. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24,* 979–992.

Potter, M. C., Dell'Acqua, R., Pesciarelli, F., Job, R., & Peressotti, F. (2005). Bidirectional semantic priming in the attentional blink. *Psychonomic Bulletin & Review, 12*(3), 460–465.

Potter, M. C., Staub, A., & O'Conner, D. H. (2002). The time course of competition for attention: Attention is initially labile. *Journal of Experimental Psychology: Human Perception and Performance, 28,* 1149–1162.

Potts, G. F. (2004). An ERP index of task relevance evaluation of visual stimuli. *Brain and Cognition, 56,* 5–31.

Raymond, J. E. (2003). New objects, not new features, trigger the attentional blink. *Psychological Science, 14,* 54–59.

Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink? *Journal of Experimental Psychology: Human Perception and Performance, 18,* 849–860.

Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1995). Similarity determines the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance, 21,* 653–662.

Reeves, A., & Sperling, G. (1986). Attention gating in short-term visual memory. *Psychological Review, 93,* 180–206.

Rolke, B., Heil, M., Streb, J., & Hennighausen, E. (2001). Missed prime words within the attentional blink evoke an N400 semantic priming effect. *Psychophysiology, 38,* 165–174.

Rolls, E. T., & Tovee, M. J. (1994). Processing speed in the cerebral cortex and the neurophysiology of visual masking. *Proceedings of the Royal Society London B, 257,* 9–15.

Sahraie, A., Milders, M., & Niedeggen, M. (2001). Attention-induced motion blindness. *Vision Research, 41,* 1613–1617.

Scharlau, I., Ansorge, U., & Horstmann, G. (2006). Latency facilitation in temporal-order judgments: Time course of facilitation as a function of judgment. *Acta Psychologica, 122,* 129–159.

Seiffert, A. E., & Di Lollo, V. (1997). Low-level masking in the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance, 23,* 1061–1073.

Sergent, C., Baillet, S., & Dehaene, S. (2005). Timing of the brain events underlying access to consciousness during the attentional blink. *Nature Neuroscience, 8,* 1391–1400.

Shapiro, K. L., Arnell, K. M., & Raymond, J. E. (1997). The attentional blink. *Trends in Cognitive Sciences, 1,* 291–296.

Shapiro, K. L., Caldwell, J., & Sorensen, R. E. (1997). Personal names and the attentional blink: A visual "cocktail party" effect. *Journal of Experimental Psychology: Human Perception and Performance, 23,* 504–514.

Shapiro, K. L., Driver, J., Ward, R., & Sorensen, R. E. (1997). Priming from the attentional blink: A failure to extract visual tokens but not visual types. *Psychological Science, 8,* 95–100.

Shapiro, K. L., & Raymond, J. E. (1994). Temporal allocation of visual attention. Inhibition or interference? In D. Dagenbach & T. H. Carr

(Eds.), *Inhibitory processes in attention, memory, and language* (pp. 151–188). San Diego, CA: Academic Press.

Shapiro, K. L., Raymond, J. E., & Arnell, K. M. (1994). Attention to visual pattern information produces the attentional blink in RSVP. *Journal of Experimental Psychology: Human Perception and Performance, 20,* 357–371.

Shapiro, K., Schmitz, F., Martens, S., Hommel, B., & Schnitzler, A. (2006). Resource sharing in the attentional blink. *NeuroReport, 17,* 163–166.

Shih, S.-I. (2008). The attention cascade model and the attentional blink. *Cognitive Psychology, 56,* 210–236.

Shih, S.-I., & Sperling, G. (2002). Measuring and modeling the trajectory of visual spatial attention. *Psychological Review, 109,* 260–305.

Shimozaki, S. S., Chen, K. Y., Abbey, C. K., & Eckstein, M. P. (2007). The temporal dynamics of selective attention of the periphery as measured by classification images. *Journal of Vision, 7,* 1–20.

Shore, D. I., McLaughlin, E. N., & Klein, R. (2001). Modulation of the attentional blink by differential resource allocation. *Canadian Journal of Experimental Psychology, 55*(4), 318–324.

Shulman, G. L., Remington, R. W., & McLean, J. P. (1979). Moving attention through visual space. *Journal of Experimental Psychology: Human Perception and Performance, 5,* 522–526.

Slagter, H. A., Lutz, A., Greischar, L. L., Francis, A. D., Nieuwenhuis, S., Davis, J., et al. (2007). Mental training affects distribution of limited brain resources. *Plos Biology, 5,* 1228–1235.

Sperling, G., & Reeves, A. (1980). Measuring the reaction time of an unobservable response: A shift of visual attention. In R. S. Nickerson (Ed.), *Attention and performance VIII* (pp. 347–360). New York: Academic Press.

Sperling, G., & Weichselgartner, E. (1995). Episodic theory of the dynamics of spatial attention. *Psychological Review, 102,* 503–532.

Suzuki, S., & Cavanagh, P. (1997). Focused attention distorts visual space: An attentional repulsion effect. *Journal of Experimental Psychology: Human Perception and Performance, 23,* 443–463.

Tomita, H., Ohbayashi, M., Nakahara, K., Hasegawa, I., & Miyashita, Y. (1999, October 14). Top-down signal from prefrontal cortex in executive control of memory retrieval. *Nature, 401,* 699–701.

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

Ungerleider, L. G., & Mishkin, M. (1982). Two cortical visual systems. In D. J. Ingle, M. A. Goodale, & R. J. W. Mansfield (Eds.), *Analysis of visual behavior* (pp. S49–S86). Cambridge, MA: MIT Press.

Visser, T. A. W. (2007). Masking T1 difficulty: Processing time and the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance, 33,* 285–297.

Visser, T. A. W., Bischof, W. F., & Di Lollo, V. (1999). Attentional switching in spatial and nonspatial domains: Evidence from the attentional blink. *Psychological Bulletin, 125,* 458–469.

Visser, T. A. W., Bischof, W. F., & Di Lollo, V. (2004). Rapid serial visual distraction: Task-irrelevant items can produce an attentional blink. *Perception & Psychophysics, 66*(8), 1418–1432.

Visser, T. A. W., Merikle, P. M., & Di Lollo, V. (2005). Priming in the attentional blink: Perception without awareness. *Visual Cognition, 12,* 1362–1375.

Visser, T. A. W., Zuvic, S. M., Bischof, W. F., & Di Lollo, V. (1999). The attentional blink with targets in different spatial locations. *Psychonomic Bulletin & Review, 6*(3), 432–436.

Vogel, E. K., & Luck, S. J. (2002). Delayed working memory consolidation during the attentional blink. *Psychonomic Bulletin & Review, 9,* 739–743.

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

Vogel, E. K., McCollough, A. W., & Machizawa, M. G. (2005, November 24). Neural measures reveal individual differences in controlling access to working memory. *Nature, 438,* 500–503.

Ward, R., Duncan, J., & Shapiro, K. (1997). Effects of similarity, difficulty, and nontarget presentation on the time course of visual attention. *Perception & Psychophysics, 59,* 593–600.

Watson, D. G., & Humphreys, G. W. (1995). Attention capture by contour onsets and offsets: No special role for onsets. *Perception & Psychophysics, 57,* 583–597.

Weichselgartner, E., & Sperling, G. (1987, November 6). Dynamics of automatic and controlled visual attention. *Science, 238,* 778–780.

Welford, A. T. (1952). The "psychological refractory period" and the timing of high-speed performance: A review and theory. *British Journal of Psychology, 43,* 2–19.

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

Wyble, B., Bowman, H., & Nieuwenstein, M. R. (in press). The attentional blink provides episodic distinctiveness: Sparing at a cost. *Journal of Experimental Psychology: Human Perception and Performance.*

Zeki, S. M. (1978, August 3). Functional specialization in the visual cortex of the rhesus monkey. *Nature, 274,* 423–428.

Zipser, K., Lamme, V. A., & Schiller, P. H. (1996). Contextual modulation in primary visual cortex. *Journal of Neuroscience, 15,* 7376–7389.

Zuvic, S. M., Visser, T. A. W., & Di Lollo, V. (2000). Direct estimates of processing delays in the attentional blink. *Psychological Research, 63,* 192–198.

(*Appendix follows*)

Appendix

Technical Details of the Model

The model consists of two parts, namely (a) sensory activity and (b) a gating system that governs entrance to working memory (WM). Both are defined by mathematical formulas that describe the activity of sets of interacting model neurons. Low-level sensory processes that lead up to the relevant perceptual representations (i.e., identity, category, and color) are not explicitly modeled, and neither are the high-level consolidation and response-generation processes in WM. Such processes, of course, take place, but within our model, they play no crucial role in explaining the data. The crucial role is played by the gating between the sensory memory and WM stages. All simulations were run in discrete time steps of 1 ms. The code used in the simulations can be found at http://olivers.cogpsy.nl.

Sensory Activity

The stream of inputs to the model are taken from the paradigms that are simulated. Of each item in the stream, the presentation duration, identity, category, and color are coded.

A first set of equations models the activity of sensory neurons that code for the identity, category (letters, digits, or symbols), and color of an item. We collectively refer to these neurons as first-stage neurons. The total activity of these sensory neurons is the product of bottom-up input and of attention already directed at the stream location, as generated by preceding items (later, we explain how activity is also convolved with attention generated by item j itself). Consider an item j that is presented from time steps t_{on} to t_{off} at location l. Sensory activity for item j at time step t, $a_{j,t}$, is the product of the bottom-up input generated by j and b_j, and of the attention directed at the location generated by preceding items, $g_{l,t}$, at a delay of 25 ms:

$$a_{j,t} = g_{l,t-25} \, b_{j,t}. \tag{A1}$$

Here, $g_{l,t-25}$ can be regarded as the "prior attentional state," as generated by all the previous items (i.e., prior to the attention yet to be triggered by the current item). The 25 ms simulates signaling delays in the effects that these previous items exert on the current item. The bottom-up input $b_{j,t}$ is governed by the following equations:

$$b_{j,t} = 0 \text{ for } t < t_{on}, \tag{A2a}$$

$$b_{j,t} = s_j \left[c_{trans}\tau_{trans}(t - t_{on})e^{-\tau_{trans}(t - t_{on})} + c_{sus}\left(1 - \frac{1}{1 + \tau_{sus}(t - t_{on})} \right) \right]$$
$$\text{for } t_{on} < t < t_{off}, \tag{A2b}$$

and

$$b_{j,t} = b_{j,t_{off}} e^{-\tau_{dec}(t - t_{off})} \text{ for } t > t_{off}. \tag{A2c}$$

The bottom-up input $b_{j,t}$ is equal to 0 before t_{on} (see Equation A2a). On the time steps that j is presented, it generates a transient (trans) signal and an additional sustained (sus) signal (see Equation A2b), both multiplied by strength s_j. Following others (e.g., Busey & Loftus, 1994), we modeled the transient signal as having the shape of a gamma distribution. Its form is determined by a shape

parameter (not made explicit in Equation A2b but set to 2) and a scale parameter τ_{trans}, with value 0.04. The scaling constant of the transient signal, c_{trans}, is equal to 3. We added the sustained signal to account for activity generated by sustained stimulus presentations: Although activity seen in brain areas involved in vision shows a marked transient, it does not decay back to zero with sustained presentation (e.g., Gawne & Martin, 2002). The sustained signal rises more slowly as the item is presented, with time constant $\tau_{sus} = 1/10$, to a maximum of c_{sus} (equal to 0.5). The resulting function resembles recorded responses of neurons in the visual system (e.g., Keysers & Perrett, 2002). More sophisticated accounts of how this activation pattern emerges from mechanisms at the individual neuronal level exist but are beyond the scope of the present model (see, e.g., Huber & O'Reilly, 2003). After presentation, $b_{j,t}$ falls exponentially (see Equation A2c), with τ_{dec} determining the speed of decay (the higher τ_{dec}, the faster decay; decay depends on masking, see below). Salient visual offsets, such as a gap within a rapid serial visual presentation task (RSVP; Lawrence, 1971) stream, are also assumed to produce transient signals (see, e.g., J. Miller, 1989; Watson & Humphreys, 1995, for behavioral evidence). In the model, gaps within a stream therefore produce the same bottom-up signal as do items.

The strength of the bottom-up signals for item j, s_j in Equation A2b, incorporates forward masking and the adaptation caused by item repetitions in a stream. Forward masking is stronger the more similar two items are, and this is implemented by making s_j a function of the contrast between item j and the item that precedes it (see Equation A3). Adaptation is implemented by making s_j a function of the distance in time between the onset of the current and the onset of the last presentation of the item:

$$s_j = \frac{1}{d_{col}d_{cat}}\left(1 - \frac{\delta_{decay}}{t_{on} - t_{on,prev}} \right). \tag{A3}$$

In Equation 3, d_{col} is the color contrast. It is equal to d_{coldif} (0.9) if j has a different color than its predecessor and to 1 if color remains the same. The contrast in category, d_{cat}, is equal to d_{catdif} (0.7, unless specified otherwise) if j is from a different category than its predecessor. It is equal to $d_{catsame}$ (1, unless specified otherwise) if the category remains the same. In accordance with "repetition blindness" effects (Kanwisher, 1987), strength s_j is assumed to be lower if the item was recently presented in the stream (i.e., the difference between the onset of the current presentation, t_{on}, and the onset of the previous presentation, $t_{on,prev}$, is small) than when it was not. Adaptation resulting from the presentation of items in previous trials is assumed to be negligible, with $t_{on} - t_{on,prev}$ effectively equal to infinity. Constant δ_{decay} was set to 50. This value results in a negative s_j if the previous presentation was less than 50 ms ago, but at that speed, the presentations are assumed to fuse and effectively become a single presentation. The strength of the decay in Equation 2c, τ_{dec}, incorporates backward masking by making it dependent on the item following j: $\tau_{dec} = .05 d_{cat} \, d_{col}$. Here, d_{col} and d_{cat} have the same values as above but are set by the contrast between item j and the following item instead

of the preceding item. This implies that τ_{dec} is smaller (and, thus, decay is slower and masking is weaker) the more dissimilar item j and the following item are.

Attentional Gating

Items generate a top-down gating response. Targets generate excitatory attention to their location, distractors generate inhibition. Both excitation and inhibition are transient signals, governed by a function with the same shape as the one for visual transients. The assumed mechanism is one where sensory activity drives gate neurons, which, in turn, modulate the sensory activity. The activation of gate neurons has effects that take time to develop and time to subside. Behaviorally, this transient response corresponds to *transient attention* (Nakayama & Mackeben, 1989). The strength of the excitatory or inhibitory gating response is assumed to depend on the strength of the bottom-up evidence for a target or distractor. For practical purposes, the model takes the sensory activity during the first 15 ms of presentation as a measure of perceptual strength. The attentional response to item j on time step t, $f_{j,t}$, is excitatory for targets (see Equation A4a) and inhibitory for distractors (see Equation A4b). Both are set relative to a baseline of 1:

$$f_{j,t} = 1 + w_j \sum_{u = t_{on}}^{t_{on} + 15} a_{j,u}[c_{att}\tau_{att}\,(t - t_{on})e^{-\tau_{att}(t - t_{on})}]$$

$$\text{for } t > t_{on} + 15 \text{ and } j: \text{target;} \quad \text{(A4a)}$$

$$f_{j,t} = \frac{1}{1 + w_j \sum_{u = t_{on}}^{t_{on} + 15} a_{j,u}[c_{att}\tau_{att}\,(t - t_{on})e^{-\tau_{att}(t - t_{on})}]}$$

$$\text{for } t > t_{on} + 15 \text{ and } j: \text{distractor.} \quad \text{(A4b)}$$

On the first 15 time steps of presentation, the summation is from t_{on} to t, instead of to $t_{on}+15$, (i.e., from onset to the current time step, so that in these 15 ms, gating does not depend on future values). Here, constant $c_{att} = 3$, and time constant $\tau_{att} = 0.015$. Strength w_j is dependent on the target in the simulated paradigm and equals $1/(d_{catdif} \times d_{coldif})$. That is, targets can be defined by category, color, or both, and the need for gating strength depends on the relative salience of the bottom-up evidence. Gaps in the stream are assumed to be slightly disruptive and are therefore assumed to generate a slight inhibitory signal, with $w_j = 0.1$. If an item has the target color but is of the alphanumeric distractor category (so it is both target-like and distractor-like), it is treated as neutral, with $w_j = 0$. Thus, $f_{j,t}$ is the attention generated by the current item, and it modulates the activity of that item at no delay.

The net attentional gating devoted to the location of the stream, l, is the product of the attention, both excitatory and inhibitory, generated by all items preceding the current item:

$$g_{l,t} = \prod_i f_{i,t}, \quad \text{(A5)}$$

with i ranging over all items preceding item j.

Entry Into WM

Within the model, entry into WM is a stochastic process. The likelihood that item j enters WM is a function of the product of the sensory activity generated by the item over the course of the stream ($a_{j,t}$, which is, in turn, a function of all prior attention; see Equation A1), of the new attention generated by the current item ($f_{j,t}$; see Equation A4), and of the likelihood that there is space in WM (i.e., that all available slots have not yet been occupied by preceding items):

$$p(j \text{ in WM}) = \left(1 - \frac{1}{1 + c_{WM}\left[\sum_t f_{j,t}a_{j,t} - r\right]^+}\right) p(q < C). \quad \text{(A6)}$$

The first part of Equation A6 is a squashing function that keeps the likelihood of entry into WM between 0 and 1. Scaling constant c_{WM} is set to 15, threshold constant r to 1. Negative values within the $[\,]^+$ brackets are set to 0. This means that if the sum of sensory activity multiplied by attention generated by item j is lower than r, then j has 0 likelihood of entering WM. $P(q < C)$ is the likelihood that the number of items already in WM, q, is lower than the capacity of WM, C. This is calculated by computing the likelihood that of the previous items in the stream, C have already entered WM. Capacity C was set to 5, which is within the range of 4 to 7 found by others (Cowan, 2001; G. A. Miller, 1956).

Simulated performance on targets is equal to the likelihood that these have entered WM at the end of the stream—$p(j$ in WM) (see Equation A6)—multiplied by .95 to allow for 5% nonspecific errors (e.g., eye blinks, general lapses of attention, wrong key presses), plus a guessing correction depending on the number of response alternatives (n) in the stream:

$$p(\text{correct}) = 0.95\,p(j \text{ in WM}) + 1/n[1 - 0.95p(j \text{ in WM})]. \quad \text{(A7)}$$

Switching Tasks and Locations

Many RSVP tasks involve switches, in which the targets are defined differently halfway through the stream (task switch) or are presented at a different location (location switch). We model such switches as stochastic. If item i is a cue to switch, the switch can occur during presentation of any subsequent item. The likelihood that this has occurred before an item j following i, $p_{switch,j}$, follows a logistic distribution:

$$p_{switch,t} = \frac{1}{1 + e^{-(t_{on,j} - t_{on,i} - \alpha)/\beta}}. \quad \text{(A8)}$$

Here, $t_{on,i}$ is the onset time of the cue item, and $t_{on,j}$ is the onset time of item j. In the case of a task switch, item j is defined as target or distractor following the new task with likelihood $p_{switch,j}$, and following the old task with likelihood $1 - p_{switch,j}$. In the case of a location switch, item j is processed with likelihood $p_{switch,j}$, whereas only a blank is processed with likelihood $1 - p_{switch,j}$. Parameter α, the mean of the distribution, is equal to 200 ms for endogenous task and locations switches and to 75 ms for exogenous cues (salient cues that automatically attract attention to their location). Parameter β, which determines the variance, was always equal to 0.2α.