

The Attentional Blink Reveals Serial Working Memory Encoding: Evidence from Virtual and Human Event-related Potentials

Patrick Craston¹, Brad Wyble², Srivas Chennu¹, and Howard Bowman¹

Abstract

■ Observers often miss a second target (T2) if it follows an identified first target item (T1) within half a second in rapid serial visual presentation (RSVP), a finding termed the attentional blink. If two targets are presented in immediate succession, however, accuracy is excellent (Lag 1 sparing). The resource sharing hypothesis proposes a dynamic distribution of resources over a time span of up to 600 msec during the attentional blink. In contrast, the ST² model argues that working memory encoding is serial during the attentional blink and that, due to joint consolidation, Lag 1 is the only case where resources are shared. Experiment 1 investigates the P3 ERP component evoked by targets in RSVP. The results suggest that, in this context, P3 am-

plitude is an indication of bottom-up strength rather than a measure of cognitive resource allocation. Experiment 2, employing a two-target paradigm, suggests that T1 consolidation is not affected by the presentation of T2 during the attentional blink. However, if targets are presented in immediate succession (Lag 1 sparing), they are jointly encoded into working memory. We use the ST² model's neural network implementation, which replicates a range of behavioral results related to the attentional blink, to generate "virtual ERPs" by summing across activation traces. We compare virtual to human ERPs and show how the results suggest a serial nature of working memory encoding as implied by the ST² model. ■

INTRODUCTION

In daily life, humans have to cope with an environment consisting of simultaneously occurring events and concurrent sensory input. In order to survive in this parallel world, attention allows us to filter out irrelevant information. On the one hand, attention lets us focus on one task at a time, whereas on the other hand, we are also often able to perform multiple tasks simultaneously. Thus, it seems that cognitive resources can be shared between tasks, suggesting a notion of divided attention. This distribution of attention, however, seems to come with concomitant costs and limitations both in terms of performance accuracy and reaction times. In this article, we investigate the extent to which attentional resources can be shared over time and the cost associated with it.

In spatial attention, the visual system was long assumed to operate in a serial manner, in that it was restricted to selecting information from only one location at a time. Attention was considered to move through the visual field in the form of a single spotlight (Eriksen & Yeh, 1985; Posner, 1980; Broadbent, 1958; von Helmholtz, 1867). However, Pylyshyn and Storm (1988), among others, disproved these classical theories by showing that humans are capable of simultaneously tracking multiple objects in

space. Some of the new theories preserve the idea of a single focus of attention, which sequentially switches between targets (Oksama & Hyöna, 2004; Pylyshyn & Storm, 1988); others propose a notion of concurrent multifocal attention, which can be focused on more than one location at a time (McMains & Somers, 2004; Awh & Pashler, 2000; Castiello & Umiltà, 1992).

Whether attention is a single spotlight, switching rapidly between locations, or whether attentional resources are distributed across multiple locations, simultaneous perception of multiple objects in space requires some notion of resource sharing. In line with this argument, Cavanagh and Alvarez (2005) conclude that the "trade-off between capacity and feature encoding (Oksama & Hyöna, 2004; Bahrami, 2003; Saiki, 2003) suggests that attention has a fixed total bandwidth for selection and the bandwidth can be shared across several input channels or targets." Hence, although the system is capable of tracking multiple objects at a time, there is a fixed amount of attentional resource. As this resource is shared across increasing numbers of targets, overall performance at the task decreases.

Recently, it has been proposed that the notion of a shared attentional resource with fixed capacity could be extended to the temporal domain (Shapiro, Schmitz, Martens, Hommel, & Schnitzler, 2006). Accordingly, if multiple target items are presented at the same spatial

¹University of Kent, UK, ²Massachusetts Institute of Technology

location within a very short period, the system allocates a certain amount of the resource to each of the targets and they are, at least to some extent, processed in a concurrent manner. Hence, if one of the targets is processed more extensively, less resource is available for other targets, which has a detrimental effect on target detection accuracy, thus, explaining a finding termed the attentional blink (AB; Raymond, Shapiro, & Arnell, 1992).

The Attentional Blink

The AB describes the empirical observation that detection of a second target (T2) is severely impaired if it follows an identified first target (T1) within less than 600 msec. This alleged blink of the “mind’s eye” (Raymond et al., 1992) was initially thought to reflect a fundamental limitation of visual perception in humans. Subsequent research, however, suggests the AB is by no means absolute. Even during the deepest part of the AB, a number of T2s are detected, that is, performance is never at zero (Raymond et al., 1992). Furthermore, evidence from electrophysiology (Rolke, Heil, Streb, & Hennighausen, 2001; Vogel, Luck, & Shapiro, 1998) and priming studies (Martens, Wolters, & van Raamsdonk, 2002; Shapiro, Driver, Ward, & Sorensen, 1997) suggests that, even if a T2 is missed during the AB, it is still processed to a semantic level. Intriguingly, if T2 is presented in immediate succession to T1, T2 detection accuracy is often above the baseline performance outside the AB (“Lag 1 sparing”; Chun & Potter, 1995).

Resource Sharing vs. Two-stage Theories

The resource sharing hypothesis suggests that the AB is an artifact of compromised allocation of attention (Shapiro et al., 2006). If the system allocates less resource to T1, more attention is available for T2 and T2 is more likely to be detected. If, however, too much resource is allocated to T1, T2 is more likely to be missed, which results in an AB (Kranzloch, Debener, Maye, & Engel, 2007).

In contrast, two-stage theories (Chun & Potter, 1995), such as the Simultaneous Type Serial Token (ST²) model (Bowman & Wyble, 2007), propose that the AB reveals a cognitive mechanism, which ensures serial working memory encoding to protect the integrity of an attentional episode (Wyble, Bowman, & Nieuwenstein, in press). If T2 is presented during the AB window, its working memory consolidation is delayed until T1 has been successfully encoded (Chun & Potter, 1995). At Lag 1, however, this “protection mechanism” breaks down and T1 and T2 are encoded into a single attentional episode (Wyble et al., in press; Chun & Potter, 1995). Joint encoding increases T2 accuracy at Lag 1, but comes at the cost of increased swaps (i.e., T1 and T2

are identified correctly but reported in the wrong order) and reduced T1 accuracy (Bowman & Wyble, 2007; Hommel & Akyürek, 2005; Potter, Staub, & O’Connor, 2002).

The P3 Component as a Measure of Resource Allocation?

A method for contrasting these theoretical positions is EEG, a noninvasive technique of recording brain activity from electrodes placed on the participant’s scalp. Event-related potentials (ERPs) are generated by averaging over segments of EEG activity time-locked to an externally generated event. The averaging process increases the observable signal by removing ongoing non-time-locked EEG activity, which is treated as background noise. The resulting ERP waveform contains a number of positive and negative deflections, which are referred to as ERP components. The P3 component of the ERP occurs 300–600 msec poststimulus presentation and is evoked most strongly by a rare event among a sequence of frequent items, through so called oddball tasks. In the context of rapid serial visual presentation (RSVP), the P3 has been argued to be a correlate of working memory update (Vogel et al., 1998).

The resource sharing hypothesis was formulated in response to a number of findings derived from EEG (Kranzloch et al., 2007; Martens, Elmallah, London, & Johnson, 2006) and MEG (Shapiro et al., 2006) experiments investigating the AB. These authors base their argument on the assumption that the size of the P3 component evoked by a target in RSVP reflects the amount of resources invested into processing this target.

However, in an extensive review of the P3 component, Kok (2001) comes to the conclusion that “the sensitivity of P3 amplitude as a measure of processing capacity has only been convincingly demonstrated in a restricted number of studies in which capacity allocation was under voluntary control, and the structural characteristics of the task (e.g., task complexity, perceptual quality of the stimuli) did not change.” Accordingly, P3 size increases if observers know beforehand that the task is going to be harder, and allocate more cognitive resource to it (Kramer & Hahn, 1995; Sirevaag, Kramer, Coles, & Donchin, 1989; Wickens, Kramer, Vanasse, & Donchin, 1983). When task difficulty is determined only by intrinsic stimulus properties, however, there is a reciprocal relationship between increasing task difficulty and P3 amplitude (Johnson, 1986).

This distinction is critical when using the P3 component to evaluate theories of the AB. During the AB, target items are often letters presented in a stream of digit distractors (e.g., Chun & Potter, 1995). Due to their shape, some target letters are masked more strongly by the distractors than others, thus, target letters can be categorized by their individual accuracy scores, yielding a measure of task difficulty according to intrinsic stimu-

lus properties. We will use the terms “easy” and “hard” to categorize letters according to their individual accuracy scores. In RSVP, target letters commonly appear in random order so that observers cannot predict whether an upcoming target in RSVP will be easy or hard, and they do not know beforehand how much resource to allocate to the target. Hence, in Kok’s (2001) terms, resource allocation is not “under voluntary control,” whereas the “structural characteristics” of the stimuli do change, and thus, the P3 should not serve as a measure of resource allocation.

Recent articles arguing in favor of resource sharing have proposed that the allocation of resource to targets in RSVP might be random, varying from trial to trial (Kranczioch et al., 2007; Shapiro et al., 2006). If, by chance, more resource is allocated to T1, less attention is available for T2, thus suggesting a tradeoff in accuracy and P3 sizes. Depending on how one interprets this argument of random allocation of resources, we can make two predictions about the resulting nature of P3 for easy and hard targets: (a) If resource allocation is truly random, it should produce no difference in the average P3 amplitude between easy and hard targets. (b) Alternatively, if hard targets are somehow able to instantaneously attract more resources, we should expect to observe a larger P3 for intrinsically hard targets, when compared to easy ones.

The ST² model, in contrast, makes a different prediction regarding the effects of target difficulty, in that the amplitude of the P3 for targets in RSVP should be mainly modulated by bottom-up strength. If a target is more easy to perceive due to its intrinsic stimulus characteristics, for instance, if it is less strongly masked, the target has more bottom-up target strength, which leads to a larger P3. Vice versa, a target that is intrinsically harder to detect will have less bottom-up strength, thus evoking a smaller P3 component.

Overview

In this article, we evaluate the resource sharing theory and the ST² model as two competing explanations of the AB. Bowman and Wyble (2007) showed that the ST² model’s neural network implementation (neural-ST²) replicates a wide range of behavioral results related to the AB. Here, we use neural-ST² to generate artificial “electrophysiological” traces by summing across activation potentials from the neural network and compare these so-called virtual ERPs (vERPs) to human ERP data (hERPs).

To this end, we first address the question of understanding P3 amplitude differences for RSVP targets, which is critical for interpreting EEG/MEG results. Does a large P3 indicate that more effort was dedicated to the task because it was harder? Or is P3 size mainly modulated by intrinsic stimulus characteristics, in which case a larger P3 indicates that the target was particularly strong,

hence, easy to perceive? This question is addressed in Experiment 1 in which participants had their EEG recorded while detecting a single letter target among digit distractors. In Experiment 1, whether a target letter is easy (or hard) depends solely on intrinsic stimulus characteristics, and thus, the hERP data (and corresponding vERPs from the ST² model) can be used to evaluate the competing hypotheses of P3 amplitude described in the previous section.

With this finding in hand, a second EEG experiment is presented in which we asked subjects to perform a two-target AB task while recording their EEG. Once again, the ST² model is used to generate corresponding vERPs. Although the resource sharing theory lacks a clear formal description, it does make a key prediction for EEG/MEG data. Resource sharing suggests that targets indirectly compete for resources during the AB through the amount of resources allocated to each of them. Hence, it predicts that the T1 P3 component should be larger for trials in which T2 is missed during the AB, as too much resource was invested into the processing of T1. On the other hand, if T2 is seen during the AB, the T1 P3 is likely to be smaller as subjects were able to allocate resources more evenly between targets. In contrast, the ST² model proposes that targets are encoded one at a time, thus emphasizing the serial nature of working memory encoding during the AB. This suggests the following prediction for the EEG/MEG correlates of target encoding during the AB. T1 consolidation (as exemplified by T1’s P3 component) should influence T2 processing in both behavioral and electrophysiological terms because T2s have to “wait” until T1 has been consolidated. The reverse, however, is not the case, that is, T1’s P3 should be unaffected, regardless of whether T2 is seen or missed, thus the influence between T1 and T2 is unidirectional. Only if the targets appear in immediate succession, as is the case at Lag 1, can there be mutual interference.

METHODS

The following section commences with a description of the ST² model and its connectionist implementation, termed “neural ST².” We explain how artificial “electrophysiological” traces (vERPs) are generated from the neural network model. Following the discussion of computational modeling methodology, we describe the methods employed in the two EEG experiments that are presented in this article.

The ST² Model

We first describe the fundamental principles of how the ST² model describes working memory, temporal attention and, in particular, the AB (for a more detailed description, please refer to Bowman & Wyble, 2007).

Types and Tokens

The ST² model employs a types–tokens account (Chun, 1997; Kanwisher, 1987) to describe the process of working memory encoding. Types describe all feature-related properties associated with an item. These include sensory properties, such as visual features (e.g., its shape, color, and the line segments comprising it) and also semantic attributes, such as a letter’s position in the alphabet. A token, on the other hand, represents episodic information, which is specific to a particular occurrence of an item, thus providing a notion of serial order. An item is encoded into working memory by creating a connection between a type and a token. At retrieval, tokens contain information about when an item occurred and, from tokens, types can be regenerated, yielding a description of what each item was and in what temporal order they appeared.

Model Architecture

As illustrated in Figure 1, the ST² model can be divided into three parts. We describe them in turn:

1. Input and extraction of types in Stage 1: Input values, which simulate target letters and digit distractors

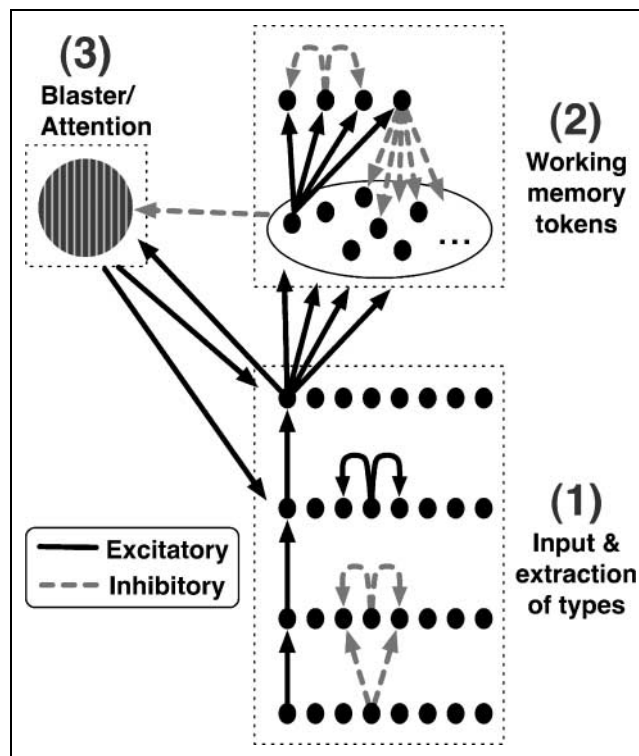


Figure 1. The ST² model: (1) Input and extraction of types in Stage 1; (2) Working memory tokens in Stage 2; (3) Temporal attention from the blaster. For the simplicity of presentation, this article contains a more abstract description of the model (please refer to Bowman & Wyble, 2007, for an extensive description of the individual layers of the ST² model).

in the RSVP stream, are fed into the model at the lowest layer of Stage 1. As activation values propagate upward, the following layers reflect forward and backward masking in early visual processing and extraction of semantic representations. A task demand mechanism operates at the highest layer of Stage 1, thus ensuring that only targets are selected for working memory encoding. Despite the fact that stimuli are presented serially during RSVP, processing within Stage 1 may exceed the presentation time of sequentially presented items. Hence, these layers are parallel or simultaneous in nature, in that more than one node can be active at any one time.

2. Working memory encoding in Stage 2: An item is encoded into working memory by connecting its type from Stage 1 to a working memory token from Stage 2. This process is referred to as “tokenization.” If at the end of a trial, the type node of a target has a valid connection to a token, the target is successfully “reported” by the ST² model. Inhibition between working memory tokens ensures that only one tokenization process is active at any one time, thus enforcing a serial nature of working memory encoding.

3. Temporal attention from the blaster: Temporal attention is implemented by a mechanism termed the blaster. Salient items in Stage 1 trigger the blaster, which provides a powerful enhancement to all nodes in the higher layers of Stage 1. The enhancement from the blaster allows targets to become sufficiently active to initiate tokenization. During tokenization, the blaster is suppressed until encoding of the target has completed. The suppression prevents a second target from re-firing the blaster while the first target is being tokenized, which would corrupt the working memory encoding process.

During the AB, T1 is being tokenized when T2 is presented, thus the blaster cannot enhance T2 as it is being suppressed by T1 tokenization. By the time T1 tokenization has completed, T2 will often lack sufficient activation to initiate its own tokenization process, which causes T2 to be missed, resulting in an AB. At Lag 1, however, T2 is presented during the window of blaster enhancement triggered by T1. T2 gains sufficient activation to join T1’s working memory encoding process and T1 and T2 are tokenized together.

Changes to the ST² Model in Comparison to Bowman and Wyble (2007)

For this work we generate vERPs from the ST² model with as few parameter changes as possible compared to the previously published version of the ST² model. Table 1 contains a list of neural network weight values that were modified. Note that we can still reproduce all behavioral data published in Bowman and Wyble (2007). The number of distractor nodes in Stage 1 is increased from 10 to 15 nodes. This has no effect on behavioral

Table 1. List of Weights that Were Modified during this Work

Layer 1 \Rightarrow Layer 2	Weight Value
Highest layer Stage 1 (“task filtered layer”) \Rightarrow Blaster	0.02003 (0.018)
Blaster recurrent excitation	0.0112 (0.01)
<i>100 msec presentation rate</i>	
First layer Stage 1 (“input layer”) \Rightarrow Second layer Stage 1 (“masking layer”)	0.023 (0.022)
<i>50 msec presentation rate</i>	
First layer Stage 1 (“input layer”) \Rightarrow Second layer Stage 1 (“masking layer”)	0.058 (0.05)

The original values (from Bowman & Wyble, 2007) are shown in brackets. All other weight values remained unchanged.

accuracy, but is required to generate 50 msec SOA vERP traces as otherwise, due to the fast presentation rate, nodes are not able to decay to baseline before being reactivated.

Each item “presented” to the model has a certain strength value. Distractors have a constant value of 0.526, whereas strength values for T1 and T2 iterate from 0.442 to 0.61 in steps of 0.014. This results in the model simulating 169 target strength combinations.

Virtual ERPs

Computational modeling is commonly focused on the replication of behavioral data. In this article, we explore an additional dimension, namely, modeling ERPs recorded during the AB. Due to the novelty of this approach, there is no established methodology for generating vERPs. Throughout this work, our philosophy is to use the most straightforward method while keeping our approach as close as possible to the mechanisms that are assumed to occur in the brain. It is obvious, however, that vERPs remain a coarse approximation of hERPs. Some factors that influence hERPs, such as the distortion of the signal by the scalp, are not addressed. Due to these limitations, one can realistically only expect to obtain a qualitative rather than a quantitative match to the data. Nevertheless, vERPs from the ST² model seem to allow us to make sensible predictions about the ERPs measured from the human scalp.

Neural Correlates of Human ERPs

The difference in electric charge between the dendrite and the postsynaptic cell body of an active neuron creates an electric dipole. To generate a signal that is strong enough to be registered by the EEG, a population of neurons has to be active together and spatially aligned, which causes the individual dipoles to summate. Cortical

pyramidal neurons have long-range connections and are aligned perpendicular to the cortex, which is why these neurons are assumed to be a major contributor of the human EEG (Luck, 2005). Pyramidal neurons release glutamate as their neurotransmitter and are therefore primarily excitatory.

Virtual ERP Calculation

The nodes in neural-ST² are organized in layers (Figure 1), which are connected via weighted connections. We assume these connections to be the analogue of synaptic projections in the brain.

vERPs from neural-ST² are generated by summing across excitatory postsynaptic node potentials, which are calculated by multiplying the activation value of a node by the weight value of the connection with the subsequent layer. We adopt the most straightforward approach and sum over all nodes of a given subset of layers in order to avoid a specific weighting of layers or normalization setting. Neurophysiological evidence suggests that there is a processing delay of around 70 msec for activation related to visual processing to travel from the retina to occipital areas (Schmolesky et al., 1998). To account for this delay, vERPs are shifted by (the model equivalent of) 70 msec.

Virtual P3

The P3 component of the hERP is commonly recorded from parietal electrode sites and considered to be a correlate of working memory encoding (Vogel et al., 1998). In the ST² model, working memory encoding occurs by creating a binding link between types from Stage 1 and tokens from Stage 2. Hence, the virtual P3 component (vP3, an example trace is shown in Figure 2) contains activation from later parts of the first stage, the nodes in Stage 2, and the binding link connecting the two stages.

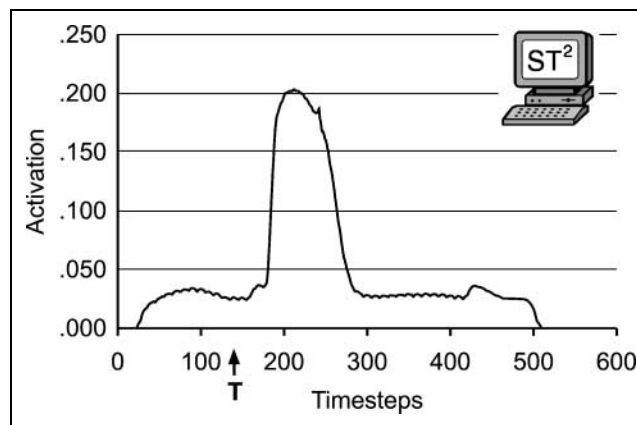


Figure 2. Virtual P3 component generated from later parts of Stage 1, Stage 2, and the binding link between the two stages. Timepoint 0 corresponds to the onset of the RSVP stream. “T” indicates the presentation of the target.

Experiment 1

Participants

Twelve university students (mean age = 24.1 years; $SD = 2.9$; 6 women; 11 right-handed) provided written consent and received 10 GBP for participation. One participant was excluded due to an excessive number of EEG artifacts, leaving 11 participants for the behavioral and EEG analysis (mean age = 24.3 years; $SD = 3.0$; 5 women; 10 right-handed). Participants were free from neurological disorders and had normal or corrected-to-normal vision. The study was approved by the local ethics committee.

Stimuli and Apparatus

We presented alphanumeric characters in black on a white background at a distance of 100 cm on a 21-in. CRT computer screen (1024×768 at 85 Hz) using the Psychophysics toolbox (Brainard, 1997) running on Matlab 6.5 under Microsoft Windows XP. Stimuli were in Arial font and had an average size of 2.1×3.4 visual angle. A photodiode verified exact stimulus presentation timing.

Procedure

Participants viewed three blocks, each consisting of 96 single target trials and four distractor-only trials. The first block was preceded by five practice trials, which were not included in the analysis. The target for each trial was chosen at random from a list of 14 capital letters (B, C, D, E, F, G, J, K, L, P, R, T, U, V); distractors could be any digit except “1” or “0.” The target’s position in the stream varied between positions 10 and 54. The “distractor only” trials were randomly inserted to make the occurrence of the target less predictable. A fixation cross presented for 500 msec preceded the first item of each stream. Items were presented at the unconventionally fast rate of approximately 20 items per second (item duration = 47.1 msec; no interstimulus interval) to ensure that participants’ detection accuracy was not at ceiling in this relatively easy single target detection task. An RSVP stream consisted of 70 items to allow a sufficient amount of time between target presentation and the end of the stream. This was required in order to prevent the subject’s behavioral response from interfering with the EEG signal evoked by the target. Each stream ended with a dot or a comma presented for 47.1 msec. Following stream presentation, participants were asked, “Was the final item a comma or a dot?” and in the following screen, “If you saw a letter, type it. If not, press Space.” Participants entered their responses using a computer keyboard. The dot–comma task was included to ensure that participants maintained their attention on the stream after the target had passed.

EEG Recording

EEG activity was recorded from Ag/AgCl electrodes mounted on an electrode cap (FMS, Munich, Germany) using a high-input impedance amplifier (1000 M Ω ; BrainProducts, Munich, Germany) with a 22-bit analog-to-digital converter. Electrode impedance was reduced to less than 25 k Ω before data acquisition (Ferree, Luu, Russell, & Tucker, 2001). EEG amplifier and electrodes employed actiShield technology (BrainProducts) for noise and artifact reduction.

The sampling rate was 2000 Hz (digitally reduced to 1000 Hz at a later stage) and the data were digitally filtered at low-pass 85 Hz and high-pass 0.5 Hz at recording. Electrodes were placed at 20 standard locations according to the international 10–20 system (Jasper, 1958). Electrooculographic (EOG) activity was bipolarly recorded from below and to the right side of the right eye. Activity from the Pz (midline parietal) electrode was used to analyze the P3 component. Because only seen targets evoke a P3, whereas missed targets do not (e.g., Kranczioch, Debener, & Engel, 2003), ERPs were generated only from trials in which the target was correctly identified.

EEG Data Analysis

EEG data were analyzed using BrainVision Analyzer (BrainProducts). The data was referenced to a common average on-line and re-referenced to linked earlobes off-line. Left mastoid acted as ground. Signal deviations in the EOG channel of more than 50 μ V within an interval of 100 msec were identified as eye blink and movement artifacts. These were removed by rejecting data in the window of 200 msec before and after an eye artifact. To verify that these trials were accurately identified by the algorithm, we performed a manual inspection after the algorithm had been applied. After artifact rejection, ERPs in each of the conditions (“easy” and “hard”) contained 531 and 387 epochs, respectively. In total, 25% of trials had to be excluded due to artifacts. ERPs were time-locked to the onset of the target and extracted from -200 to 1200 msec with respect to target onset. After segmentation, direct current drift artifacts were removed using a DC detrend procedure employing the average activity of the first and last 100 msec of a segment as starting and end point, respectively. Following this, the baseline was corrected to the prestimulus interval (-200 msec to timepoint 0) and segments were averaged to create ERPs. Unless otherwise stated, ERP component amplitudes were derived from mean amplitude values within a certain window. ERP component latencies were calculated using 50% area latency analysis (Luck & Hillyard, 1990). Amplitude and latency values from subject averages were submitted to Matlab scripts (Trujillo-Ortiz, Hernandez-Walls, Castro-Perez, & Barba-Rojo, 2006; Trujillo-Ortiz, Hernandez-Walls, &

Trujillo-Perez, 2004) to perform repeated measures ANOVA. Where appropriate, *p* values were adjusted using Greenhouse–Geisser correction. After all statistical analyzes, a 25-Hz low-pass filter was applied to enhance visualization of ERP components.

Computational Modeling

In order to simulate single-target RSVP streams with 50 msec presentation rate, the input patterns presented to the model contained 40 items with the target appearing at position 14 of the stream. Each item was presented for 10 timesteps, the equivalent of 50 msec.

Experiment 2

Participants

Twenty new under- and postgraduate university students (mean age = 23.1 years, *SD* = 3.2; 10 women; 18 right-handed) provided written consent and received 10 GBP for participation. Two participants were excluded from the analysis. The first one seemed to be a nonblinker (Martens, Munneke, Smid, & Johnson, 2006), as his performance was at ceiling across all three lags. The second participant was excluded due to persistently high oscillations in the alpha band throughout the experiment. Hence, 18 participants remained for behavioral and EEG analysis (mean age = 22.5 years, *SD* = 2.7; 9 women; 18 right-handed). Participants were free from neurological disorders and had normal or corrected-to-normal vision. The study was approved by the local ethics committee.

Stimuli and Apparatus


Stimulus presentation was equal to that in Experiment 1 except for a reduction in average stimulus size ($1.03^\circ \times 0.69^\circ$ visual angle) to ensure that the paradigm produced a reliable AB effect.

Procedure

Participants viewed four blocks of 100 trials. Before starting the experiment, participants were asked to make five eye blinks and five horizontal eye movements to configure the algorithm for eye blink artifact rejection. Participants performed eight practice trials, which were not included in the analysis. As shown in Figure 3, RSVP streams were preceded by a fixation cross in the center of the screen. After 400 msec, the cross turned into an arrow indicating the side at which the targets would be presented. After 200 msec, two streams of digits were simultaneously presented at an equal distance of 2.6° visual angle to the left and right of fixation.¹ The RSVP stream consisted of 35 items presented for 105.9 msec each with no interstimulus interval. For 84% of trials in a block, the stream on the side indicated by the arrow contained two targets (T1 and T2), in 16% of trials both streams were made up of distractor digits only. The “distractor only” trials were randomly inserted to make the occurrence of targets less predictable. In a trial, T1 and T2 were selected from a list of 14 possible targets (A, B, C, D, E, F, G, H, J, K, L, N, P, R, T, U, V, Y); distractors could be any digit except “1” or “0.” T1 appeared between positions 5 and 17; T2 followed T1 at position 1 (no intervening distractors—Lag 1), position 3 (2 intervening distractors—Lag 3) or position 8 (7 intervening distractors—Lag 8). The arrow remained in the center of the screen until the streams were over and then turned into either a dot or a comma.

Before the experiment started, participants were told to keep their eyes fixated on the center of the screen, as trials with eye movements would have to be excluded. Participants were told to direct their covert attention toward the indicated stream, search for the two target letters, and remember whether the last item was a dot or a comma. Participants were informed that streams could contain either two or zero targets. Following stream presentation, participants were presented with the message “If you saw letters—type them in order, then dot or

Figure 3. Dual-target bilateral RSVP stream.



comma for the final item” and entered their response without time pressure using a computer keyboard.

EEG Recording and Data Analysis

EEG methods for Experiment 2 were the same as for Experiment 1, with the following changes. The sampling rate was 1000 Hz and the data were filtered at 80 Hz low-pass and 0.25 Hz high-pass at recording. Horizontal eye movements, recorded from a bipolar EOG channel placed below and to the left of the participant’s left eye, indicated that participants had moved their eyes away from fixation and toward one of the RSVP streams. These trials, along with trials violating the artifact rejection procedure described for Experiment 1, were excluded from further analysis. In total, 10% of trials had to be excluded due to artifacts. After artifact rejection, ERPs for the conditions contained the following number of trials: Lag 3 noAB—863 epochs; Lag 3 AB—702 epochs; Lag 8—1201 epochs; Lag 1—946 epochs.² For Experiment 2, ERPs were time-locked to T1 and extracted from –200 to 1800 msec with respect to T1 onset.

Computational Modeling

In order to simulate two-target RSVP streams with 100 msec presentation rate, the input patterns presented to the model were composed of 25 items presented for 20 timesteps (equivalent to 100 msec) each. T1 appeared at position 7 in the RSVP stream and T2 followed T1 with 0 to 7 distractors (Lags 1–8) between the two targets.

RESULTS

Experiment 1

Behavior

We determine the accuracy score for each target letter by using the behavioral results for T1 accuracy per letter from a previously published study (Bowman & Wyble, 2007), which employed a similar RSVP paradigm.⁵ Accordingly, all targets are classified as belonging either to the “easy” or the “hard” group of target letters. By dividing targets a priori (with respect to the experiment reported here), we counter arguments that our subdivision into easy and hard reflects random variation in attentional state (i.e., alertness) of subjects, rather than fluctuations in intrinsic stimulus strength. The fact that it is the same letters that are easy (respectively hard) in the Bowman and Wyble (2007) experiment and the experiment reported here is strong evidence that variation in intrinsic stimulus characteristics underlies this subdivision.

The behavioral results from Experiment 1 of this study show that the “hard” target letters (E, C, B, P, F, J, and R) have an average accuracy of 67% ($SEM = 4$), whereas the “easy” target letters (T, K, U, V, L, D, and G) have an

average accuracy of 86% ($SEM = 3$). The difference in accuracy scores between the easy and the hard target groups is highly significant [$F(1, 10) = 48.26$, $MSE < .01$, $p < .001$].

In the ST^2 model, a target is classified as hard if its strength value is less than or equal to the value of distractors (strength values ranging from 0.442 to 0.526). Target values above those of distractors contribute to the easy condition (strength values ranging from 0.540 to 0.610). The ST^2 model provides a qualitative fit of the behavioral accuracy scores for the hard (ST^2 accuracy: 57%) and easy (ST^2 accuracy: 100%) conditions.

Human ERP

As seen in Figure 4, the P3 for easy targets has a significantly larger amplitude than the P3 for hard targets [$F(1, 10) = 9.65$, $MSE = 2.1$, $p = .011$]. The mean amplitude in the 300–600 msec posttarget area is 9.7 μV ($SEM = 1.1$) for easy targets and 7.8 μV ($SEM = 1.4$) for




Figure 4. (A) hERP P3 component from Pz for the easy and hard conditions. Positive is plotted upward. (B) ST^2 's vERP containing the virtual P3 component for the easy and hard conditions. For both panels, “T” indicates the presentation of the target and ERPs are time-locked to presentation of the target.

hard targets. Although the P3 for hard targets starts slightly later than the P3 for easy targets, it also returns back to baseline more rapidly, thus, the small difference in 50% area latency analysis (Luck & Hillyard, 1990) is nonsignificant [easy targets: mean = 464 msec ($SEM = 8$) vs. hard targets: mean = 469 msec ($SEM = 13$); $F(1, 10) = 0.55$, $MSE = 241$, $p = .476$].

Virtual ERP

In the ST^2 model, easy targets have higher input strength, and thus, generate more activation than hard targets. Figure 4 illustrates how the vP3 is larger in amplitude for easy compared to hard targets (mean vP3 amplitude: Easy 0.203 vs. Hard 0.189). Once target activation reaches later parts of Stage 1, easy targets trigger an earlier blaster response, which causes these items to be encoded into working memory more rapidly. The result is a slightly earlier vP3 component for easy (vP3 50% area latency: 455 msec equivalent) compared to hard (vP3 50% area latency: 460 msec equivalent) targets, as seen in Figure 4.

Experiment 2

Behavior

Attentional Blink. As shown in Figure 5, human accuracy at identifying T2 (conditional on correct report of T1) shows a significant effect of lag [$F(2, 17) = 15.58$, $MSE = 0.03$, Greenhouse–Geisser (GG)- $\epsilon = .74$, $p < .001$]. Pairwise comparisons emphasize the presence of an AB. T2 accuracy is significantly lower at Lag 3 compared to Lag 8 [$F(1, 17) = 11.66$, $MSE = 0.03$, $p = .003$] and Lag 1 [$F(1, 17) = 60.88$, $MSE = 0.01$, $p < .001$]. If T2 is presented in immediate succession to T1 (Lag 1), T2 accuracy is significantly higher than T2 accuracy at Lag 8 [$F(1, 17) = 5.41$, $MSE = 0.01$, $p = .033$]. The ST^2 model replicates a U-shaped AB curve. T2 accuracy (conditional on correct report of T1) is reduced at Lag 3 compared to Lag 8 and Lag 1. Furthermore, T2 accuracy at Lag 1 is slightly higher than at Lag 8 (Figure 5).

When comparing performance of the ST^2 model to these data, it should be noted that the model was configured to replicate a specific set of AB data (Chun & Potter, 1995). Subsequent studies (including this experiment) mostly reported higher Lag 3 accuracy, thus, a less drastic AB effect. To keep with the philosophy of changing as few parameters as possible compared to the ST^2 model published in Bowman and Wyble (2007), we sacrifice a perfect quantitative fit of the data from this experiment and, instead, emphasize the replication of a qualitative AB effect.

Reduced T1 Accuracy at Lag 1. Observers are significantly worse at reporting T1 if T2 is presented at Lag 1 compared to when T2 is presented at Lag 3 [$F(1, 17) =$

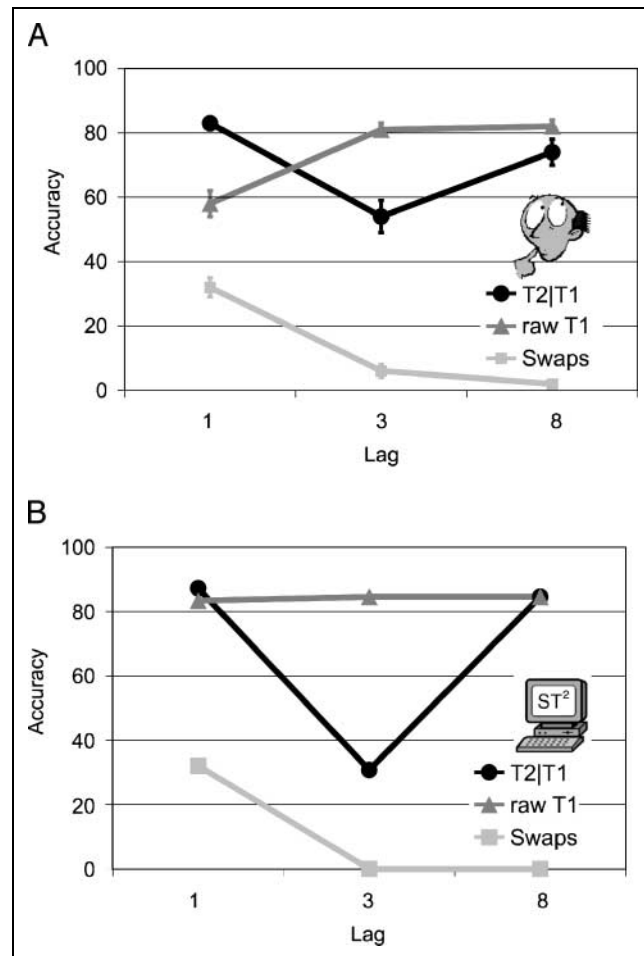


Figure 5. (A) Human behavioral accuracy data for Lag 1, Lag 3, and Lag 8. (B) Simulated behavioral accuracy of the ST^2 model for Lag 1, Lag 3, and Lag 8. Circles indicate T2 accuracy conditional on correct T1 report, triangles represent raw T1 accuracy, and squares indicate swaps, that is, the condition when T1 and T2 were correctly identified but reported in the wrong order.

49.68, $MSE = 0.01$, $p < .001$] or Lag 8 [$F(1, 17) = 61.21$, $MSE = 0.01$, $p < .001$]. The ST^2 model replicates a reduction in T1 accuracy at Lag 1.

No Effect on T1 Accuracy When T2 is at Lag 3 or 8. We observe no significant difference in T1 accuracy between T2 being presented at Lag 3 or Lag 8 [$F(1, 17) = 0.44$, $MSE < 0.01$, $p = .515$; Figure 5]. Furthermore, there is no difference in T1 accuracy whether an AB occurs or not [T1 accuracy conditional on seen T2 at Lag 3: 79%, $SEM = 4$; T1 accuracy conditional on missed T2 at Lag 3: 78%, $SEM = 3$; $F(1, 17) = 0.03$, $MSE = 0.02$, $p = .862$]. The ST^2 model replicates these effects, as simulated T1 accuracy is at baseline irrespective of whether T2 is presented at Lag 3 or Lag 8.

Increased Number of Swaps at Lag 1. At Lag 1 we observe a high percentage of swaps, but swaps are negligible at Lag 3 and 8. The difference in swaps between

Lag 1 and Lag 3 [$F(1,17) = 58.67, MSE = 0.01, p < .001$] and also Lag 1 compared to Lag 8 [$F(1, 17) = 133.31, MSE = 0.01, p < .001$] is highly significant. The ST² model replicates this effect and produces a high proportion of swaps if T2 is presented at Lag 1 but produces no order inversions at Lags 3 and 8.

Human ERPs

Our results suggest no significant difference in mean amplitude of T1's P3 (300–600 msec) with respect to T2 presentation (Figure 6). First, there is no significant difference in T1 P3 amplitude whether an AB occurs or not [Lag 3 AB: 6.5 μ V ($SEM = 0.6$) vs. Lag 3 noAB: 7.3 μ V ($SEM = 0.6$); $F(1, 17) = 1.91, MSE = 2.7, p = .185$]. Second, there is no significant difference in T1 P3 amplitude

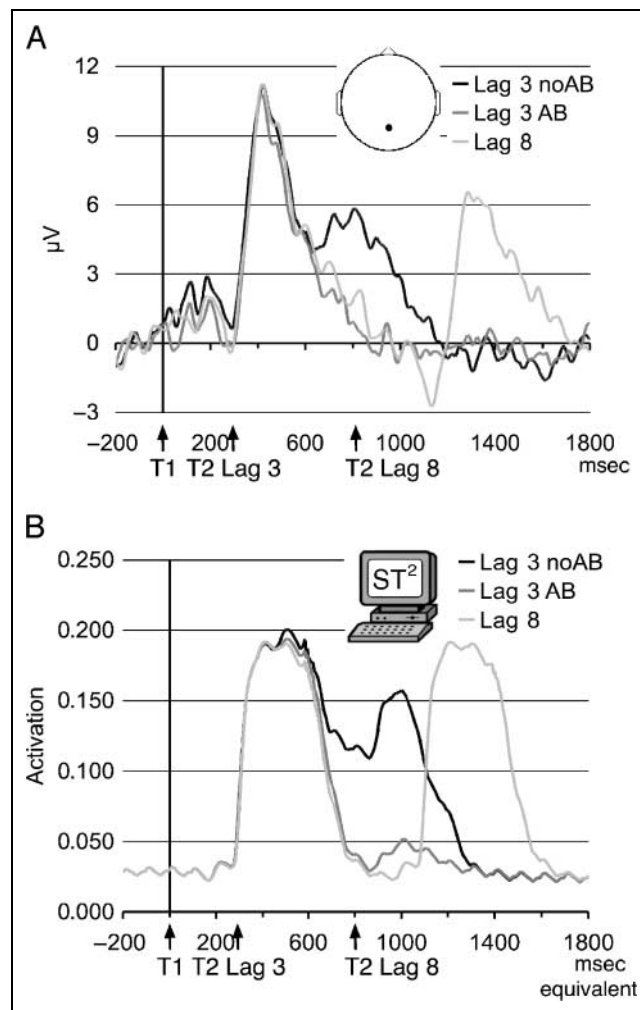


Figure 6. (A) hERP from Pz for (1) the Lag 3 noAB condition (T1 and T2 correctly reported), (2) the Lag 3 AB condition (T1 accurately identified but T2 not correctly reported), and (3) the Lag 8 condition (T1 and T2 correctly reported). Positive is plotted upward. (B) ST²'s vERP containing the virtual P3 for (1) the Lag 3 noAB condition, (2) the Lag 3 AB condition, and (3) the Lag 8 condition. In both panels, "T1" and "T2" indicate the presentation of T1 and T2, respectively, and ERPs are time-locked to T1.

whether T2 is presented at Lag 3 or Lag 8 [Lag 3 noAB: 7.3 μ V ($SEM = 0.6$) vs. Lag 8: 7.0 μ V ($SEM = 0.6$); $F(1, 17) = 0.32, MSE = 2.0, p = .576$].

As suggested by Figure 6, T1 P3 50% area latency (calculated for the 300–600 msec window) seems to be independent of T2 presentation. First, there is no significant difference in T1 P3 latency whether an AB occurs or not [Lag 3 AB: 453 msec ($SEM = 5$) vs. Lag 3 noAB: 452 msec ($SEM = 5$); $F(1, 17) = 0.02, MSE = 241.8, p = .883$]. Second, whether T2 is presented at Lag 3 or Lag 8 has no significant effect on T1 P3 latency [Lag 3 noAB: 452 msec ($SEM = 5$) vs. Lag 8: 454 msec ($SEM = 3$); $F(1, 17) = 0.18, MSE = 191.9, p = .670$].

We replicate the finding that T2 evokes a P3 component in those trials in which an AB does not occur (Figure 6; see also Kranczioch et al., 2003). The difference in mean amplitude in the 600–1200 msec window between the AB and noAB condition is highly significant [Lag 3 AB: 0.7 μ V ($SEM = 0.6$) vs. Lag 3 noAB: 3.4 μ V ($SEM = 0.6$); $F(1, 17) = 24.58, MSE = 2.6, p < .001$].

Figure 7 suggests the presence of a joint P3 for T1 and T2 if T2 is presented at Lag 1. The mean P3 amplitude in the 300–600 msec window is significantly larger than the mean amplitude for the same window if T2 is presented at Lag 8 [Lag 1: 8.5 μ V ($SEM = 0.5$) vs. Lag 8: 7.0 μ V ($SEM = 0.6$); $F(1, 17) = 11.03, MSE = 1.77, p = .004$].

Virtual ERPs

According to the ST² model, at Lag 3 and Lag 8 targets are encoded into working memory in a serial fashion. If T2 is presented at Lag 3, the blaster is suppressed by T1's encoding process and T2's tokenization is delayed. However, a T2 presented at Lag 8 appears after T1 has been encoded into working memory, thus, the T2 can initiate a new encoding process.

As shown in Figure 6, there is no difference in the mean amplitude of T1's vP3 amplitude, irrespective of whether or not an AB occurs at Lag 3 or whether T2 is presented at Lag 8 (Lag 3 noAB: 0.18; Lag 3 AB: 0.18; Lag 8: 0.18). There is also no difference in 50% area latency for T1's vP3 component between the Lag 3 AB, the Lag 3 noAB condition and the Lag 8 condition (Lag 3 AB: 470 msec equivalent; Lag 3 noAB: 470 msec equivalent; Lag 8 470 msec equivalent). In line with serial working memory encoding, at Lag 3 and Lag 8 T2 is presented beyond the timepoint where it could have an effect on T1's tokenization.

T2 items that are presented at Lag 3 and have relatively low target strength are not encoded into working memory. They show only a small deviation from baseline in the vERP (Figure 6; T2 vP3 mean amplitude for Lag 3 AB: 0.06), which remains below threshold. T2s that are strong enough to "outlive" T1's tokenization, however, re-fire the blaster once T1 encoding has completed. They are consolidated into working memory




Figure 7. (A) hERP from Pz for (1) the Lag 1 condition (T1 and T2 correctly reported), (2) the Lag 8 condition (T1 and T2 correctly reported). Positive is plotted upward. (B) ST²'s vERP containing the virtual P3 component for (1) the Lag 1 condition and (2) the Lag 8 condition. In both panels, "T1" and "T2" indicate the presentation of T1 and T2, respectively, and ERPs are time-locked to T1.

and show a vP3 component (Figure 6; T2 vP3 mean amplitude for Lag 3 noAB: 0.13).

According to the ST² model, the targets are jointly encoded into working memory at Lag 1. T2 is presented within the period of T1's blaster enhancement and joins into T1's tokenization process. Hence, the vERP in Figure 7 contains one joint vP3 component for both T1 and T2 at Lag 1. The joint vP3 at Lag 1 combines bottom-up activation of two targets, which is reflected in a larger area under the vP3 curve for the Lag 1 vP3 compared to a vP3 for an individual target, that is, T1's vP3 if T2 is presented at Lag 8 (Lag 1: 0.28 vs. Lag 8: 0.17).

DISCUSSION

The present study addresses two issues central to the evaluation of theories of the AB using electrophysiol-

ogy. In Experiment 1, we investigate the effect of task difficulty on the P3 component evoked by a target presented in RSVP. Various hypotheses provide conflicting predictions on the relationship between task difficulty and the P3, in that if the target is harder to detect, the amplitude of the P3 should (a) increase (Martens, Elmallah, et al., 2006), (b) remain equal (Kranczioch et al., 2007; Shapiro et al., 2006), or (c) decrease (Kok, 2001). The second experiment does not find a modulation of T1 processing by T2 presented during the AB, thus, our data are in contrast with previously published findings (Kranczioch et al., 2007; Martens, Elmallah, et al., 2006; Shapiro et al., 2006). We evaluate the findings from our Experiment 2 and then discuss the discrepancy between our data and the previous experimental findings.

The Meaning of P3 Amplitude for Targets in RSVP

The results from Experiment 1 provide evidence in favor of the P3 component for targets in RSVP being a correlate of bottom-up target strength. First, certain target letters have significantly higher accuracy scores than others. We use the behavioral data from a previous study (Bowman & Wyble, 2007) to classify target letters as being easy or hard. Our results replicate the previous finding and show a highly significant difference in accuracy between easy and hard letters. This suggests that there are consistent differences in target strengths, which are determined by the identity of each target letter. Such a measure of task difficulty is purely due to intrinsic stimulus characteristics. As target letters are presented at random, observers cannot predict whether a target is going to be easy or hard.

Second, the P3 amplitude is significantly larger for easy compared to hard targets. This finding contradicts theories based on the assumption that P3 amplitude reflects the amount of resource allocated to processing a target in RSVP. According to such theories, more resource should be required to process harder targets (Martens, Elmallah, et al., 2006). In consequence, we should find a larger P3 for hard targets, however, the data from Experiment 1 shows the opposite effect. Alternatively, P3 size might be determined by the amount of resource allocated to the processing of the target, which more or less randomly fluctuates from trial to trial (Kranczioch et al., 2007; Shapiro et al., 2006). However, this hypothesis predicts that a measure of task difficulty due to intrinsic stimulus characteristics (as employed in Experiment 1) should not modulate P3 amplitude, which is in contrast with our results. Hence, based on the results of Experiment 1, we can conclude that if pre-allocated effort is either random or equal in every trial, as can be assumed due to the randomness of target presentation in RSVP, intrinsic target strength is a main modulator of P3 amplitude.

In neural network terms, target strength might be referred to as bottom-up trace strength. One of the main

arguments in the theory underlying the ST^2 model is that the working memory encoding process is influenced by the target's strength. A stronger target will be consolidated into working memory in a more durable manner, which is reflected in a larger vP3 component. Hence, the findings from Experiment 1 validate and support the ST^2 model.

Working Memory Encoding is Serial during the Attentional Blink

Both the ST^2 model and the resource sharing theory propose that T1 processing affects the consolidation of T2 during the AB, which is supported by behavioral (e.g., Chun & Potter, 1995) and EEG (Vogel et al., 1998) data. In addition to the unidirectional influence of T1 on T2, however, resource sharing also argues that there is mutual interference during the AB, as T1 and T2 compete indirectly through the amount of resource allocated to them. The behavioral and EEG data from Experiment 2, however, do not support this hypothesis. These data suggest that T2 does not influence T1 if presented at Lag 3 or Lag 8. In addition, there is no effect on T1 processing whether an AB occurs or not.

Our findings support theories that suggest T1 and T2 do not compete for resources during the AB (Olivers, 2007) and are consistent with the hypothesis of serial working memory encoding during the AB (Bowman & Wyble, 2007). If T2 is presented at Lag 3, T1 is in the process of being encoded into working memory. During T1's tokenization process, the attentional enhancement is suppressed, preventing any interference from T2. Providing T2 has sufficient activation strength, T2's working memory encoding process is delayed until T1 has been consolidated. If T2, however, is too weak, it is lost and an AB occurs.

The data from Experiment 2 are thus in contrast with a key prediction from the resource sharing theory. However, resource sharing—as it stands—lacks a formal interpretation, leaving open the possibility of uncertainty over the exact predictions of the theory. One might thus imagine a modified version of the theory, which would explain the data presented in this article, while, nevertheless, remaining within the “umbrella” of resource sharing. In that eventuality, however, the resource sharing theory risks becoming “unfalsifiable.”

Interference between T1 and T2 at Lag 1

If T1 and T2 are presented in immediate succession (i.e., at Lag 1), the serial mechanism of working memory encoding is not enforced. As indicated by the results from Experiment 2, T1 and T2 seem to be encoded into working memory together, thus evoking a single P3 component. This finding is a replication of the MEG results reported in Kessler et al. (2005), who report a single

M300 component for T1 and T2 at Lag 1. The increase in swaps at Lag 1 provides evidence for joint consolidation during Lag 1 sparing, which sometimes leads to a loss of order information for T1 and T2 (Bowman & Wyble, 2007). With respect to the shape of the P3 component at Lag 1, neither the human nor the virtual P3 components appear to consist of two individual P3s for T1 and T2 that are offset by 100 msec. As the P3 is larger in amplitude but not much broader in time, this suggests a single P3 component (indicating a single enhanced encoding process) for two target items, which is in line with the theory proposed by the ST^2 model.

As long as target characteristics are relatively simple (single letters), the joint consolidation has a beneficial effect on T2 accuracy, as exemplified by the Lag 1 sparing effect (Bowman & Wyble, 2007). There is a negative effect on T1 accuracy, however, as it is reduced if T2 is presented at Lag 1 (see Figure 5 and also Hommel & Akyürek, 2005).

Hence, if there exists some aspect of resource sharing in time, it occurs if targets are presented in immediate succession, as is the case at Lag 1. According to the ST^2 model, T1 receives an attentional enhancement from the blaster, which lasts for around 150 msec. As long as T2 is presented within this period, T2 can join the encoding process and resources are shared between the two targets.

Spreading the Sparing

If more than two targets are presented in a row, however, a number of studies have shown that subjects are capable of reporting these targets without showing an AB (Olivers, van der Stigchel, & Hulleman, 2007; Nieuwenstein & Potter, 2006; Di Lollo, Kawahara, Ghorashi, & Enns, 2005). It seems as if Lag 1 sparing can be extended to more than two targets and longer time periods (see also Bowman, Wyble, Chennu, and Craston, 2008). The current version of the ST^2 model cannot account for spreading the sparing; for work on a revised version of the model that addresses these findings please refer to Wyble et al. (in press).

Evaluating Previous Findings

As previously mentioned, a number of recent articles investigating the AB using EEG (and MEG) techniques have argued in favor of resource sharing during the AB. The data from those studies seem to be in contrast with this article's findings and predictions from the ST^2 model. In the following section, we take a closer look at these previous results. The data presented in each of the articles in question are tested against the following set of criteria, which we believe an EEG/MEG experiment should fulfill in order to provide evidence for resource sharing during the AB.

P3 as a Measure of Resource Allocation?

Demonstrate that the size of the P3 component evoked by a target in RSVP can be used as a measure of the cognitive resource/effort invested into the detection of that target.

Resource Sharing during the AB?

Resource sharing proposes that if more cognitive resources are allocated to T1, the T2 is more likely to be missed. Accordingly, the P3 component for T1 should be larger for those trials in which an AB occurs compared to when T2 is detected and there is no AB.

McArthur, Budd, and Michie (1999)

This study investigates the relationship between T1-related processing (as exemplified by its P3 component) and the AB. Both the P3 component and the AB are “maximal at about 300 msec” and return to baseline around 700 msec following the presentation of T1, thus, it seems that “the AB and P300 [or P3] follow a similar time course” (McArthur et al., 1999).⁴ Indeed, a significant correlation between the amplitude of six time intervals of the T1 P3 (235–325 msec, 328–415 msec, 415–505 msec, 505–595 msec, 595–685 msec, 685–775 msec; grand-averaged across all lags of T2 presentation) and the depth of the AB⁵ at Lags 1–6 (Figure 2 in McArthur et al., 1999) emphasizes the similarity between the time course of T1’s P3 and the AB.

P3 as a Measure of Resource Allocation?

In McArthur et al. (1999), difficulty is not manipulated on the basis of intrinsic stimulus characteristics (as in Experiment 1 of this article) but by making T1 less or more frequent. The authors assume that frequent targets are easy and infrequent targets are hard to perceive. However, the data from Martens, Elmallah, et al. (2006, p. 209) suggests the opposite, that is, lower average accuracy scores for frequent than infrequent targets, although the results are not significant (p values of approximately .10). Consequently, the relationship between frequency and task difficulty in the AB context is unclear.

Furthermore, due to the very nature of P3, the less frequent a target is, the more of an “oddball” it becomes (Kok, 2001). Thus, P3 size is likely to be strongly modulated by frequency/oddball effects, which may not be related to the difficulty of identifying the stimulus, or to the amount of resources allocated to it. With this point in mind, the finding of less frequent targets eliciting a larger P3 (Figure 4 in McArthur et al., 1999) does not per se provide evidence for the P3 component as a measure of resource allocation and does not contradict our results from Experiment 1.

Resource Sharing during the AB?

As T1 P3 data for the Lag 3 noAB condition are not presented in McArthur et al. (1999), this study cannot directly contribute toward the current discussion. However, McArthur et al. find a negative correlation between T1 P3 size and depth of the AB ($r = -.59$, $p = .03$; Figure 3 in McArthur et al., 1999), which provides evidence against resource sharing but in favor of a reciprocal relationship during the AB (Bowman et al., 2008).

Martens, Elmallah, et al. (2006)

This article investigates cueing and frequency effects on the AB. In Experiment 1, T1 difficulty is modulated by making T1 more or less frequent. In Experiment 2, T1 difficulty is manipulated by presenting a cue (the same letter as the T1) above the RSVP stream shortly before the presentation of T1.

P3 as a Measure of Resource Allocation?

Experiment 1 in Martens, Elmallah, et al. (2006) is a replication of McArthur et al. (1999) in that a notion of task difficulty is modified by making T1 more or less frequent. As discussed in the previous section, we argue that the relationship between task difficulty and frequency is unclear. What is clear is that frequency alone is a potent factor in determining P3 size (Kok, 2001), which explains a larger P3 for infrequent targets than for frequent targets (Figure 1 in Martens, Elmallah, et al., 2006) without resorting to the explanations involving task difficulty or resource allocation.

We believe that the results from Experiment 2 can be explained by the way in which T1 was cued. Cueing the T1 with the same character makes it easier to detect in behavioral terms, however, also makes the T1 less of an oddball, which explains the decrease in P3 amplitude for targets preceded by valid cues compared to invalid cues (Figure 3 in Martens, Elmallah, et al., 2006). Furthermore, invalidly cued T1s also come as more of a “surprise” to the participant, which increases the amplitude of the P3 component (Kok, 2001; Donchin, 1981). Hence, these results per se do not provide evidence in favor of the P3 being a measure of resource allocation as they are confounded by frequency and expectancy effects influencing P3 amplitude.

Resource Sharing during the AB?

Both experiments presented in Martens, Elmallah, et al. (2006) show T1’s P3 to be smaller⁶ on those trials in which no AB occurs compared to when T2 is missed and AB does occur, thus suggesting resource sharing. However, if T1’s P3 is mainly modulated by frequency and expectancy effects, as suggested in the previous paragraph, the data support a different conclusion. By increasing

the frequency of T1 or by validly cueing it, the AB is attenuated (Tables 1 and 2 in Martens, Elmallah, et al., 2006), which is in line with the reciprocal relationship between T1 strength and the AB (Bowman et al., 2008). Hence, the noAB condition is likely to contain a larger number of frequent T1s (Experiment 1) and validly cued T1s (Experiment 2) than the AB condition. Smaller T1 P3s in the noAB compared to the AB condition (Figures 2 and 4 in Martens, Elmallah, et al., 2006) can be explained by the reduction of T1 P3 amplitude through increased frequency and valid cueing effects. Hence, we argue that the differences in P3 size between the noAB and the AB condition do not per se support resource sharing during the AB.

As it stands, further investigation is needed to provide evidence for resource sharing. Such a study would manipulate task difficulty using intrinsic stimulus characteristics, in order to avoid experimental confounds from various factors affecting P3 size.

Shapiro et al. (2006)

This study presents M300 (MEG P3 equivalent) data for both T1 and T2 during the AB. Task difficulty is not manipulated, hence, cannot be discussed.

Resource Sharing during the AB?

The difference in T1 M300 amplitude between the AB and noAB conditions at Lag 2 is not significant ($p > 0.05$), hence, on this measure the data cannot provide evidence for resource sharing. However, the authors do find that T1 M300 amplitude is reduced if T2 is presented inside compared to outside the AB window, which suggests that T2 is able to influence T1 processing during the AB. Such a finding is in contrast with the ST² model's proposal of serial working memory encoding during the blink. A potential explanation for the finding might be the experimental setup of the study. There is evidence for interference between targets at Lag 1, so a T2 presented at Lag 2 might be presented close enough to influence T1 processing. Other studies (Experiment 2 of this study or Martens, Munneke, et al., 2006), which use Lag 3 as the AB condition, do not find a modulation of T1's P3, hence, the evidence is inconclusive.

Shapiro et al. (2006) report a positive correlation between the size of a subject's T1 M300 and the "strength" of their AB impairment. They argue that this is evidence for resource sharing, as it indicates that if a subject is able to allocate less resource to T1 (exemplified by a smaller T1 M300) they are able to reduce their AB deficit. However, such a positive correlation between T1 P3 size and depth of the AB was not found in other previously published studies (Martens, Elmallah, et al., 2006; McArthur et al., 1999).

Furthermore, we believe there might be an additional confound. What if certain participants always have smaller

M300 components (for both T1 and T2) than other participants? If, as reported for blinkers and nonblinkers (Martens, Munneke, et al., 2006), these participants are also worse at the behavioral task, that is, have a stronger AB, this would produce the positive correlation observed in Shapiro et al. (2006), emphasizing individual differences in the behavioral and MEG data. However, it requires a study showing a significant positive correlation between T1 M300 (or P3) size and the depth of the AB within each subject, for instance, across experimental blocks, to prove resource sharing.

Krancioch et al. (2007)

In this article, the authors present an EEG study of the AB including data containing the P3 component for T1 and T2. As task difficulty is not manipulated, this issue is not discussed.

Resource Sharing during the AB?

Krancioch et al. (2007) report a "significant interaction of the factors T2 performance and time window [levels T1-P3 window and T2-P3 window] [$F(1, 14) = 5.25, p = .038$]" when T2 is presented at Lag 2, that is, during the AB (see Figure 2B in Krancioch et al., 2007). They conclude that "the T1-related P3 process is larger for trials in which T2 is missed, whilst the T2-related P3 process is smaller in these trials" and that there is resource sharing during the AB.

We argue, however, that the significant interaction does not necessarily provide evidence for resource sharing. The factor time window consists of two levels, namely, "T1-P3" and "T2-P3," whereas the factor T2 performance consists of the levels "T2 seen" and "T2 missed." Although the interaction indicates a relationship between T2 performance and P3 time window, such an analysis is not necessarily evidence for a modulation of the "T1-P3" by the AB.

We illustrate this by performing an equivalent statistical analysis on our data from Experiment 2. A time window ("T1-P3" and "T2-P3") by T2 accuracy ("T2 seen" and "T2 missed") interaction analysis on our data is also significant [$F(1,17) = 7.72, MSE = 3.5, p = .0129$]. Two separate paired tests, however, indicate that the interaction is due to a highly significant relationship between T2 accuracy and "T2-P3" [$F(1, 17) = 24.58, MSE = 2.6, p < .001$], whereas a comparison of "T1-P3" and T2 accuracy is not significant [$F(1, 17) = 1.91, MSE = 2.7, p = .185$]. Hence, without a paired test between "T1-P3" and T2 accuracy, the data from Krancioch et al. (2007) do not necessarily provide evidence for resource sharing.

Martens, Munneke, et al. (2006)

This article is not directly related to the current discussion as it is primarily concerned with the difference in

EEG signatures between so-called blinkers and non-blinkers. They do, however, make an interesting observation concerning T1 P3 latency, which is relevant to the resource sharing discussion.

Resource Sharing during the AB?

Martens, Munneke, et al. (2006) report delayed T1 consolidation if T2 is presented at Lag 3 compared to Lag 8. This finding suggests that T2 can have some influence on T1 if presented at Lag 3, which is intriguing and, indeed, troublesome for the ST² model. The reported delay in T1 P3 latency for T2 inside compared to outside the AB, however, resulted from peak latency analysis [Lag 3: 495 msec, Lag 8: 427 msec, $t(10) = 2.275$, $p = .046$; S. Martens, personal communication, January 2007]. Luck (2005) suggests that if ERP components overlay in time, as is the case during the AB, a 50% area latency analysis (Luck & Hillyard, 1990) can yield more reliable results. The present study and others (Kranzioch et al., 2007; Martens, Elmallah, et al., 2006; Shapiro et al., 2006) do not find a delay in T1 consolidation if T2 is presented at Lag 3 compared to Lag 8, thus, the evidence in favor of delayed T1 consolidation during the AB is inconclusive.

Evaluating the vERP Technique

The vERPs presented here are used to validate the computational model but also provide opportunities for electrophysiological experimentation strategies. A review by Picton et al. (2000) emphasizes the importance of a clear hypothesis before conducting EEG experiments: “The overwhelming amount of ERP data along the time and scalp distribution dimensions can easily lead to incorrect post hoc conclusions based on trial-and-error analyzes of multiple time epochs and electrode sites.” Virtual ERPs provide a means of making more formal predictions of ERP latencies and amplitudes, which can aid the construction of hypotheses prior to experimental design and data collection. One can investigate how parameter changes in the model affect results in both the virtual behavioral and virtual “electrophysiological” domain, thereby giving a principled method for exploring a theoretical hypothesis.

Due to the nature of EEG, the extraction of signals related to the cognitive processes of interest from background activity can be problematic. The vERP, however, can be dissected into its underlying components. For example, one could generate vERP traces related to attentional processes or working memory consolidation by including only the associated parts of the model. If one used blind source separation techniques, such as Independent Components Analysis (Makeig, Debener, Onton, & Delorme, 2004), to decompose the hERP, correlations between individual components of the vERP and hERP

might help to further explain the cognitive processes underlying the hERP.

Conclusion

In this article, we present findings from two electrophysiological studies addressing issues fundamental to the evaluation of current theories of temporal attention and the AB. We use the ST² model and its neural network implementation to generate vERP traces, which we compare to the hERPs. In addition to validating the dynamics of the computational model, the vERPs are used to make predictions from the theory underlying the ST² model.

Experiment 1 suggests that, at least for targets in RSVP, the P3 component is modulated mainly by target strength and provides only a limited measure of the amount of resource allocated to the task. Thus, EEG/MEG experiments that were taken in support of the resource sharing theory, which assumed P3 size to be a measure of cognitive resource allocated, might have to be reinterpreted.

In Experiment 2, our data suggest that if two targets are presented in immediate succession and within a very short period (<150 msec), they can be encoded into working memory together. However, during the AB, our data suggest that the encoding of the first target into working memory influences the consolidation of subsequent targets, but this interference is not mutual. Thus, “resource sharing in time” seems to be limited to short time spans (<150 msec) and cannot be extended to the duration of the AB.

To recapitulate the issue of dividing an attentional resource among multiple tasks, we can conclude that although such a mechanism seems to exist in the spatial domain (Cavanagh & Alvarez, 2005), resource sharing in temporal attention is severely limited. When orienting in space, the system seems to be able to dynamically adapt its behavior to achieve an effective tradeoff between monitoring the visual field and looking at individual items in detail. In time, however, such dynamic adaptation is restricted to very short periods (i.e., Lag 1) where it is constrained by the length of an attentional episode. Thus, as suggested by the ST² model, the AB is an observable side effect of this strategy, which enforces a notion of serial order and ensures that perception of stimuli in time is unambiguous.

Acknowledgments

This work was supported by the UK Engineering and Physical Research Council under grant number GR/S15075/01 awarded to Howard Bowman and a Doctoral Training Account awarded to Patrick Craston. We thank the three reviewers and Dinkar Sharma for helpful comments contributing towards this work.

Reprint requests should be sent to Patrick Craston, Computing Laboratory, University of Kent, Canterbury, Kent, CT2 7NF, UK, or via e-mail: patrick@craston.de.

Notes

1. This study employed a bilateral RSVP paradigm as we also investigated a modulation of the lateralized N2pc component during the AB (Chennu, Craston, Wyble, & Bowman, 2008). Target presentation to the left and right of fixation was equally probable, randomized and the P3 was recorded from the midline Pz electrode. Hence, bilateral presentation was irrelevant for the purpose of this study.
2. In the following sections, “Lag 3 noAB” refers to the conditions when T2 was presented at Lag 3 and both targets were correctly identified, thus, an AB did not occur. “Lag 3 AB” is the condition when T1 was accurately reported but T2 could not be correctly identified, hence, the observer experienced an AB on that particular trial. The “Lag 8” and “Lag 1” conditions describe scenarios in which T2 was presented at the given lag (with respect to T1) and both targets were correctly reported.
3. The 54-msec SOA experiment from Bowman and Wyble (2007) also used a presentation rate of approximately 20 items per second and the resulting T1 accuracy (averaged across conditions where T2 is presented at Lag 12/648 msec, Lag 14/756 msec, and Lag 16/864 msec) is comparable to the accuracy of detecting single targets in the current experiment (72% vs. 77%).
4. Note that the similarity in time course of the P3 component and AB is increased by shifting the whole AB curve forward in time by 235 msec. This is justified by the need to account for “the propagation delay between probe [the T2] onset and the arrival of the signal [processing related to T1] at the cortex” (McArthur et al., 1999), in order for the T1 to be processed to a level where it could influence the processing of T2.
5. The term “depth of the attention blink” is the opposite of T2 performance, that is, how strong the AB impairment (and thus how low T2 accuracy) is at that particular timepoint.
6. Note the effect seems rather weak. In Experiment 1, statistical significance is at $p = .085/p = .048$ (peak amplitude/400–520 msec mean value) when comparing the T1 P3 of the AB to the noAB condition. In Experiment 2, significance levels are at $p = .050/p = .062$ (peak amplitude/432–584 msec mean value) when comparing the T1 P3 of the AB to the noAB condition.

REFERENCES

- Awh, E., & Pashler, H. (2000). Evidence for split attentional foci. *Journal of Experimental Psychology: Human Perception and Performance*, *26*, 834–846.
- Bahrami, B. (2003). Object property encoding and change blindness in multiple object tracking. *Visual Cognition*, *10*, 949–963.
- Bowman, H., & Wyble, B. (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.
- Brainard, D. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 433–436.
- Broadbent, D. (1958). *Perception and communication*. London: Pergamon Press.
- Castiello, U., & Umiltà, C. (1992). Splitting focal attention. *Journal of Experimental Psychology: Human Perception and Performance*, *18*, 837–848.
- Cavanagh, P., & Alvarez, G. A. (2005). Tracking multiple targets with multifocal attention. *Trends in Cognitive Sciences*, *9*, 349–354.
- Chennu, S., Craston, P., Wyble, B., & Bowman, H. (2008). Transient attentional enhancement during the attentional blink: ERP correlates of the ST² model. In R. French & E. Thomas (Eds.), *Proceedings of the 10th neural computation and psychology workshop*. London: World Scientific.
- Chun, M. (1997). 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., & Potter, M. (1995). A two-stage model for multiple target detection in rapid serial visual presentation. *Journal of Experimental Psychology: Human Perception and Performance*, *21*, 109–127.
- Di Lollo, V., Kawahara, J., Ghorashi, S., & Enns, J. (2005). The attentional blink: Resource depletion or temporary loss of control? *Psychological Research*, *69*, 191–200.
- Donchin, E. (1981). Surprise! ... Surprise? *Psychophysiology*, *18*, 493–513.
- Eriksen, C., & Yeh, Y. (1985). Allocation of attention in the visual field. *Journal of Experimental Psychology: Human Perception and Performance*, *11*, 583–597.
- Ferree, T. C., Luu, P., Russell, G. S., & Tucker, D. M. (2001). Scalp electrode impedance, infection risk, and EEG data quality. *Clinical Neurophysiology*, *112*, 536–544.
- Hommel, B., & Akyürek, E. (2005). Lag-1 sparing in the attentional blink: Benefits and costs of integrating two events into a single episode. *Quarterly Journal of Experimental Psychology A*, *58*, 1415–1433.
- Jasper, H. (1958). The ten–twenty electrode system of the International Federation. *Electroencephalography and Clinical Neurophysiology*, *10*, 371–375.
- Johnson, R. (1986). A triarchic model of P300 amplitude. *Psychophysiology*, *23*, 367–384.
- Kanwisher, N. (1987). Repetition blindness: Type recognition without token individuation. *Cognition*, *27*, 117–143.
- 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.
- Kok, A. (2001). On the utility of P3 amplitude as a measure of processing capacity. *Psychophysiology*, *38*, 557–577.
- Kramer, A. F., & Hahn, S. (1995). Splitting the beam: Distribution of attention over noncontiguous regions of the visual field. *Psychological Science*, *6*, 381–386.
- Kranczoch, C., Debener, S., & Engel, A. (2003). Event-related potential correlates of the attentional blink phenomenon. *Cognitive Brain Research*, *17*, 177–187.
- Kranczoch, C., Debener, S., Maye, A., & Engel, A. (2007). Temporal dynamics of access to consciousness in the attentional blink. *Neuroimage*, *37*, 947–955.
- Luck, S. (2005). *An introduction to the event-related potential technique*. Cambridge: MIT Press.
- Luck, S., & Hillyard, S. (1990). Electrophysiological evidence for parallel and serial processing during visual search. *Perception & Psychophysics*, *48*, 603–617.
- Makeig, S., Debener, S., Onton, J., & Delorme, A. (2004). Mining event-related brain dynamics. *Trends in Cognitive Sciences*, *8*, 204–210.
- Martens, S., Elmallah, K., London, R., & Johnson, A. (2006). Cuing and stimulus probability effects on the P3 and the AB. *Acta Psychologica*, *123*, 204–218.
- Martens, S., Munneke, J., Smid, H., & Johnson, A. (2006). Quick minds dont 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.
- McArthur, G., Budd, T., & Michie, P. (1999). The attentional blink and P300. *NeuroReport*, 10, 3691–3695.
- McMains, S., & Somers, D. (2004). Multiple spotlights of attentional selection in human visual cortex. *Neuron*, 42, 677–686.
- Nieuwenstein, M., & Potter, M. (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.
- Oksama, L., & Hyöna, J. (2004). Is multiple object tracking carried out automatically by an early vision mechanism independent of higher-order cognition? An individual difference approach. *Visual Cognition*, 11, 631–671.
- Olivers, C. (2007). The time course of attention: It is better than we thought. *Current Directions in Psychological Science*, 16, 11–15.
- Olivers, C., van der Stigchel, S., & Hulleman, J. (2007). Spreading the sparing: Against a limited-capacity account of the attentional blink. *Psychological Research*, 71, 126–139.
- Picton, T., Bentin, S., Berg, P., Donchin, E., Hillyard, S., Johnson, R., et al. (2000). Guidelines for using human event-related potentials to study cognition: Recording standards and publication criteria. *Psychophysiology*, 37, 127–152.
- Posner, M. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, 32, 3–25.
- Potter, M., Staub, A., & O'Connor, D. (2002). The time course of competition for attention: Attention is initially labile. *Journal of Experimental Psychology: Human Perception and Performance*, 28, 2.
- Pylyshyn, Z., & Storm, R. (1988). Tracking multiple independent targets: Evidence for a parallel tracking mechanism. *Spatial Vision*, 3, 179–197.
- Raymond, J., Shapiro, K., & Arnell, K. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink? *Journal of Experimental Psychology: Human Perception and Performance*, 18, 849–860.
- 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.
- Saiki, J. (2003). Feature binding in object-file representations of multiple moving items. *Journal of Vision*, 3, 6–21.
- Schmolesky, M., Wang, Y., Hanes, D., Thompson, K., Leutgeb, S., Schall, J., et al. (1998). Signal timing across the macaque visual system. *Journal of Neurophysiology*, 79, 3272–3278.
- Shapiro, K., Driver, J., Ward, R., & Sorensen, R. (1997). Priming from the attentional blink: A failure to extract visual tokens but not visual types. *Psychological Science*, 8, 95–100.
- Shapiro, K., Schmitz, F., Martens, S., Hommel, B., & Schnitzler, A. (2006). Resource sharing in the attentional blink. *NeuroReport*, 17, 163–166.
- Sirevaag, E., Kramer, A., Coles, M., & Donchin, E. (1989). Resource reciprocity: An event-related brain potentials analysis. *Acta Psychologica*, 70, 77–97.
- Trujillo-Ortiz, A., Hernandez-Walls, R., Castro-Perez, A., & Barba-Rojo, K. (2006). *epsGG: Greenhouse-Geisser epsilon. A MATLAB file*. Retrieved on June 17, 2008, from www.mathworks.com/matlabcentral/fileexchange/loadFile.do?objectId=12839.
- Trujillo-Ortiz, A., Hernandez-Walls, R., & Trujillo-Perez, R. (2004). *RMAOV1: One-way repeated measures ANOVA. A MATLAB file*. Retrieved on June 17, 2008, from www.mathworks.com/matlabcentral/fileexchange/loadFile.do?objectId=5576.
- Vogel, E., Luck, S., & Shapiro, K. (1998). Electrophysiological evidence for a postperceptual locus of suppression during the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance*, 24, 1656–1674.
- von Helmholtz, H. (1867). *Treatise on physiological optics*. Volume 3, J.P.C. Southall, transl (1962). New York: Dover.
- Wickens, C., Kramer, A., Vanasse, L., & Donchin, E. (1983). Performance of concurrent tasks: A psychophysiological analysis of the reciprocity of information-processing resources. *Science*, 221, 1080–1082.
- Wyble, B., Bowman, H., & Nieuwenstein, M. (in press). The attentional blink and episodic distinctiveness: Sparing at a cost. *Journal of Experimental Psychology: Human Perception and Performance*.