

Electrophysiological Evidence for a Postperceptual Locus of Suppression During the Attentional Blink

Edward K. Vogel and Steven J. Luck
University of Iowa

Kimron L. Shapiro
University of Wales, Bangor

When an observer detects a target in a rapid stream of visual stimuli, there is a brief period of time during which the detection of subsequent targets is impaired. In this study, event-related potentials (ERPs) were recorded from normal adult observers to determine whether this "attentional blink" reflects a suppression of perceptual processes or an impairment in postperceptual processes. No suppression was observed during the attentional blink interval for ERP components corresponding to sensory processing (the P1 and N1 components) or semantic analysis (the N400 component). However, complete suppression was observed for an ERP component that has been hypothesized to reflect the updating of working memory (the P3 component). Results indicate that the attentional blink reflects an impairment in a postperceptual stage of processing.

Over the past several decades, the vast majority of studies of visual attention have examined the operation of attention across space. In the visual search task, for example, a target item must be detected within an array of distractor items that are presented at different locations from the target. In recent years, however, there has been growing interest in the operation of attention over time rather than over space. Time-based attention has typically been studied with the rapid serial visual presentation (RSVP) task, which can be thought of as a time-based analog of the visual search task. Specifically, whereas visual search tasks typically require the observer to detect a target in an array of items presented simultaneously at different locations, RSVP tasks typically require the observer to detect a target in a sequence of items presented sequentially at a single location. Because visual search and RSVP are so analogous, it seems reasonable to suppose that the same mechanisms of attention would operate in both tasks. However, space plays a uniquely important role in the organization of the primate visual system, and several investigators have proposed that space also plays a unique role in visual attention (Hillyard &

Münste, 1984; Nissen, 1985; Treisman & Gelade, 1980; Tsal & Lavie, 1993). Consequently, special attentional mechanisms may be used for space-based attention that are not available for time-based attention. In this study we addressed this issue by comparing the nature of attentional selection in a time-based attention task with previous findings from space-based attention tasks.

Time-Based Attention and the Attentional Blink

Much of the recent interest in time-based attention derives from a phenomenon that was independently discovered by Broadbent and Broadbent (1987) and by Sperling and colleagues (Reeves & Sperling, 1986; Weichselgartner & Sperling, 1987) and that was subsequently named the *attentional blink* by Raymond, Shapiro, and Arnell (1992). The attentional blink is typically observed in RSVP tasks in which the stimuli are presented at a rate of approximately 10 items per second and the observers are required to detect or discriminate two targets on each trial. The two targets are typically called T1 and T2, and the lag between T1 and T2 is typically varied randomly across trials (e.g., on a Lag 3 trial, T2 is the third item after T1). In experiments of this nature, the correct detection of T1 is followed by a period of time during which the observers frequently fail to detect T2, and this period of impaired T2 detection is called the "attentional blink" because it is similar to the sort of impairment that would be produced if an eyeblink were triggered by T1. Specifically, T2 detection accuracy typically drops to its lowest point at a T1–T2 lag of three items and then recovers by a lag of six to eight items (Chun & Potter, 1995; Maki, Frigen, & Paulson, 1997; Raymond et al., 1992). Interestingly, there is usually little or no impairment at Lag 1, although this depends on the details of the experimental design (see, e.g., Duncan, Ward, & Shapiro, 1994).

The attentional blink is analogous to impairments of distractor processing that have been observed in visual search tasks. For example, Tsal and Lavie (1988) asked observers to identify a visual search target of a particular

Edward K. Vogel and Steven J. Luck, Department of Psychology, University of Iowa; Kimron L. Shapiro, Department of Psychology, University of Wales, Bangor, Wales.

This study was supported by Grant 95-38 from the McDonnell-Pew Program in Cognitive Neuroscience. Portions of the data in Experiment 2 were previously published in "Word Meanings Can Be Accessed but Not Reported During the Attentional Blink," by S. J. Luck, E. K. Vogel, and K. L. Shapiro, 1996, *Nature*, 382, pp. 616–618.

We thank the following people for contributing directly or indirectly to the ideas and data discussed in this article: Lyndon Berglan, Marvin Chun, Vince Di Lollo, Howard Egeth, Jim Hoffman, Bill Maki, and Jane Raymond.

Correspondence concerning this article should be addressed to Edward K. Vogel or Steven J. Luck, Department of Psychology, University of Iowa, E11 Seashore Hall, Iowa City, Iowa 52242-1407. Electronic mail may be sent to edward-vogel@uiowa.edu or steven-luck@uiowa.edu.

color and then report as many additional items from the array as possible. They found that the observers were impaired at reporting items that were far from the target location compared with items that were adjacent to the target (for related findings, see Chelazzi, Miller, Duncan, & Desimone, 1993; Kim & Cave, 1995; Luck, Fan, & Hillyard, 1993). This parallels the typical attentional blink pattern, in which T2 is reported accurately when it immediately follows T1 but is missed when it is temporally more distant. Performance eventually recovers after the attentional blink, whereas no recovery is observed at larger distances in visual search. However, this difference probably reflects the fact that attention is eventually withdrawn from T1 in the RSVP paradigm; visual search performance would similarly be expected to recover once attention was withdrawn from the target.

Although the RSVP task is analogous to the visual search task and the attentional blink phenomenon in the RSVP task is analogous to the suppression of distractor items in the visual search task, there is good reason to believe that the underlying attentional mechanisms are not entirely identical. Specifically, whereas attentional selection in spatial attention paradigms appears to begin at an early stage, before stimulus identification is complete, it appears that the attentional blink reflects a postperceptual attentional mechanism that operates at the stage of working memory. In the following sections we discuss the research that supports these proposals.

An Early Locus of Space-Based Selection

Theories of attention have generally been divided between *early-selection theories*, which propose that attention influences perceptual processes, and *late-selection theories*, which propose that attention operates only after perception is complete. The terms *early* and *late* have been construed in a variety of ways, but the most fundamental distinction is whether attentional selection occurs before or after a stimulus has been fully identified (for particularly clear statements of this distinction, see Duncan, 1980; Duncan & Humphreys, 1989). In this article, we also use the terms *perceptual* and *postperceptual* to refer to this same distinction.¹

There is little debate about whether attention can sometimes operate at a late, postperceptual stage. For example, there are many situations in which an observer will identify a stimulus but choose not to make a response to it, which is clearly a case of postperceptual selective processing. The primary focus of the locus-of-selection debate has thus been the question of whether attention sometimes operates at an early stage, before perception is complete. Despite decades of research, however, the traditional behavioral techniques of cognitive psychology have been unable to firmly determine whether early selection is possible. The primary difficulty is that these techniques rely on measurements of behavioral output, which reflects the combined effects of early and late processes. This makes it difficult to determine with certainty the stage of processing that is influenced by a given experimental manipulation. Consider, for example,

the Eriksen flanker paradigm, in which to-be-ignored distractor items are found to interfere with the processing of a to-be-attended target item (see C. W. Eriksen, 1995). When such interference is found, it is often used as evidence that the distractors are fully identified and that selection therefore occurs at a postperceptual stage (e.g., B. A. Eriksen & Eriksen, 1974; Hagenaar & van der Heijden, 1986). However, the observed interference could just as easily reflect a perceptual-level selection mechanism that only partially or probabilistically attenuates the distractors (see Treisman, 1964). Similarly, when no interference is observed, this result is often used to argue for early selection (e.g., Francolini & Egeth, 1980; Lavie, 1995; Yantis & Johnston, 1990). However, if the interference normally arises at the stage of response selection, it could easily be eliminated by an attentional mechanism that operates at a stage that is interposed between stimulus identification and response selection (see Allport, Tipper, & Chmiel, 1985; Driver & Tipper, 1989). Thus, neither the presence nor the absence of interference from distractors can unambiguously be used to determine the locus of selection.

To avoid ambiguities of this nature, several attention researchers have turned to event-related potential (ERP) recordings, which provide a continuous measure of processing between a stimulus and a response and can therefore be used to pinpoint the time at which attention begins to influence processing. ERPs are scalp-recorded voltage fluctuations that reflect synchronous neuronal activity associated with sensory, motor, or cognitive events (for reviews, see Donchin, Karis, Bashore, Coles, & Gratton, 1986; Hillyard & Picton, 1987). As shown in Figure 1, the ERP waveform is composed of several negative and positive voltage fluctuations that are triggered by an event such as a visual stimulus. The peaks and valleys of the waveform are called *waves* or *components* and reflect the temporal progression of processing, beginning with sensory activity and progressing through higher level cognitive processes and response-related processes. Each ERP component is typically named according to its polarity and temporal position within the overall ERP waveform (e.g., *P2* for the second positive component, or *P215* to indicate a precise peak

¹ This division of cognitive processing into perceptual and postperceptual stages is crude because it cannot easily accommodate factors such as top-down influences on perception and the use of partial information by decision and response processes. Similarly, given the possibility that a stimulus may be reinterpreted after further information has been acquired, there may never be a specific time at which a stimulus has been "fully identified." However, these difficulties can be minimized by defining the terms *perceptual* and *postperceptual* in a task-specific manner. Specifically, the *completion of perception* can be defined as the point at which the observer has extracted the stimulus attributes that are to be reported for the specific task being performed. According to this definition, a picture of a red spoon would be considered fully identified in a color discrimination task as soon as its color was explicitly represented, but this same picture would not be considered fully identified until its name became available in a naming task or until its membership in the silverware category was determined in a categorization task.




Figure 1. Example event-related potential (ERP) waveforms. Note that negative is plotted upward and that Time 0 represents the onset of the ERP-eliciting stimulus. A: Example of the effects of spatial attention on the P1 and N1 components. B: Example of the effects of semantic mismatch on the N400 component. C: Example of the effects of stimulus probability on the P3 component.

latency of 215 ms poststimulus). By determining which components are influenced by attention, it is possible to determine whether attention begins to operate at an early or a late stage.

Several ERP studies of visuospatial attention have demonstrated that—under appropriate conditions—attention influences processing within 100 ms of stimulus onset (for reviews, see Hillyard et al., 1996; Hillyard, Mangun, Woldorff, & Luck, 1995; Luck, 1998). As illustrated in Figure 1A, the P1 component is typically found to be larger for attended-location stimuli than for ignored-location stimuli, and this effect typically begins 70–100 ms after stimulus onset (see, e.g., Luck et al., 1994; Mangun, Hillyard, & Luck, 1993; Neville & Lawson, 1987). Notably, the same P1 attention effect is observed for both task-relevant target stimuli and for task-irrelevant nontarget stimuli that bear no resemblance to the targets (Heinze, Luck, Mangun, & Hillyard, 1990; Luck et al., 1993). The early onset latency of this effect and its insensitivity to the identity of the ERP-eliciting stimulus provide strong evidence that visuospatial attention begins to operate at an early stage, before stimulus identification is complete. In addition, a positron

emission tomography study has shown that the P1 attention effect is accompanied by a change in blood flow in the ventral occipital lobe (Heinze et al., 1994), and a single-unit recording study has demonstrated similar attentional modulations in areas V2 and V4 of monkey visual cortex (Luck, Chelazzi, Hillyard, & Desimone, 1997). Together, these studies provide strong evidence that visuospatial attention begins to operate at an early stage, at least under some conditions.

A Late Locus of Time-Based Selection?

Although there is substantial evidence indicating that space-based selection begins at an early stage, the current evidence suggests that time-based selection operates only after perception is complete, although this issue has received relatively little study. The most compelling reason to predict a late locus of selection in a time-based task such as RSVP is that there appears to be no need for selective perceptual processing in RSVP: The visual system appears to be able to identify stimuli at such a rapid rate that there is typically no need to limit perceptual processing to a subset of the incoming stimuli. For example, Potter (1976) found that untrained observers could accurately extract the main theme from pictures of complex real-world scenes presented at rates of more than 8 pictures per second. These scenes were much more complex than the simple alphanumeric characters or words typically used in RSVP tasks. It is therefore reasonable to suppose that the visual system is capable of fully identifying every item in a typical RSVP sequence. Thus, there would seem to be no reason for the use of a perceptual-level attentional mechanism in most time-based attention tasks. Potter did, however, indicate that attentional selection should be important at a later stage. Specifically, she found that observers became inaccurate at high rates of presentation if the task required them not only to perceive the pictures but also to store them in working memory. This suggests that working memory is overloaded in RSVP tasks and that attentional selection may play an important role in determining what information is encoded or consolidated in working memory.

More direct evidence for a postperceptual locus of selection in the attentional blink paradigm was provided by Shapiro, Driver, Ward, and Sorensen (1997), who used an RSVP task with three targets (T1, T2, and T3). When T2 was presented during the attentional blink period triggered by T1—and was therefore not correctly discriminated by the observers—T2 still significantly primed T3. This priming effect indicates that T2 was identified even though the observers could not accurately report it. The most plausible explanation for this finding is that the visual system fully identified every item in the RSVP stream but that T2 was not stored in a retrievable form in working memory, presumably because the working memory system was busy with T1. Because T2 was fully identified, however, it was able to cause perceptual-level priming of T3 (see also Shapiro, Caldwell, & Sorensen, 1997). Similarly, Maki et al. (1997) found that distractor stimuli presented during the attentional blink period can prime a subsequently presented T2 stimu-

lus, and this also indicates that stimuli are fully identified during the attentional blink.

Although the findings of Shapiro, Driver, et al. (1997) indicate that some information about T2 identity was extracted by the visual system during the attentional blink period, these findings do not completely rule out the possibility that the attentional blink is at least partially attributable to an early selection mechanism. Specifically, it is possible that some perceptual-level suppression was present during the attentional blink, but not enough to completely eliminate priming on every trial. This possibility could not be addressed by Shapiro, Driver, et al. because the experimental design did not permit a comparison of the magnitude of priming during the attentional blink period with the magnitude of priming outside of this period. Thus, although there are several good reasons to suspect that the attentional blink reflects a postperceptual suppression of processing, unambiguous evidence for this hypothesis has not yet been provided. The main purpose of the present study was to provide this evidence.

Overview of the Present Study

In the present study, we applied the ERP technique to the attentional blink paradigm to determine the stage at which processing is impaired during the attentional blink. More specifically, we examined the ERP waveforms elicited by stimuli presented during the blink period to establish which ERP components, and thus which processing stages, are affected by the attentional blink. The specific ERP components that were examined are illustrated in Figure 1. The first experiment examined the P1 and N1 components, which reflect perceptual processing: They have an early onset time, are sensitive to stimulus parameters such as brightness and spatial frequency, and are evoked obligatorily regardless of task. As shown in Figure 1A, these components are typically suppressed for stimuli presented at ignored locations in spatial attention experiments (Clark & Hillyard, 1996; Luck et al., 1994; Mangun & Hillyard, 1991). In the first experiment, we examined these components during an RSVP task to determine whether sensory processing would be similarly suppressed during the attentional blink.

Figure 1B illustrates the N400 component, which was examined in Experiment 2. The N400 component is highly sensitive to the degree of mismatch between a word and a previously established semantic context (Besson, Kutas, & Van Petten, 1992; Kutas & Hillyard, 1980a; Kutas, Van Petten, & Besson, 1988; Osterhout & Holcomb, 1995; Rugg, 1984). For example, a large N400 would be elicited by the last word of the sentence "He got in the car and turned on the mustard" but not by the last word of the sentence "He got in the car and turned on the radio." Because a word must be identified before its meaning can be compared with a previously established semantic context, the presence of an N400 component can be used to indicate that a word has been identified. Thus, in Experiment 2 we sought to determine whether words presented during the attentional blink period would elicit an N400 component, which would

indicate that they were identified even though they could not be reported.

In the final experiment, we examined the P3 wave. As illustrated in Figure 1C, the amplitude of the P3 wave is inversely related to the probability of the task-defined category of the ERP-eliciting stimulus (see the review by Johnson, 1986). For example, if observers are asked to categorize names as male and female, and if male names are presented more frequently than female names, then the female names will elicit larger P3 waves than will the male names. It is important that P3 amplitude is sensitive to the probability of the category, as defined by the task, rather than to the probability of the individual stimulus. Note, however, that the P3 component does not appear to reflect the categorization process per se but instead reflects a process that follows categorization. Specifically, the leading hypothesis of the P3 wave proposes that it reflects the updating of working memory (Donchin, 1981; Donchin & Coles, 1988). Because we hypothesized that the attentional blink reflects an impairment at the stage of working memory, we predicted that this component would be suppressed during the attentional blink interval.

Experiment 1

The first experiment tested the hypothesis that the attentional blink reflects a suppression of sensory processing. We tested this hypothesis by examining the effects of the attentional blink on the P1 and N1 waves, which are the first major positive and negative ERP components typically elicited by a visual stimulus. These components reflect sensory processes and are primarily sensitive to the physical characteristics of the eliciting stimulus, such as brightness (Hillyard & Picton, 1987). Numerous studies have found that these two components can be modulated by spatial attention (Clark & Hillyard, 1996; Luck et al., 1994; Mangun & Hillyard, 1991). That is, stimuli presented at ignored locations elicit smaller P1 and N1 waves than stimuli presented at attended locations. If the same mechanisms of attention are also responsible for the attentional blink, smaller P1 and N1 components should be observed for stimuli presented during the attentional blink period than for stimuli presented outside of this period. However, if the attentional blink operates at a later processing stage, no suppression of the P1 and N1 waves should be observed during the attentional blink.

The use of ERP recordings in RSVP tasks leads to some technical difficulties that place important constraints on the experimental design. In particular, each item presented in an RSVP stream produces an ERP response that lasts for several hundred milliseconds after stimulus onset, long past the onset of the next stimulus in the stream. Consequently, the ERP elicited by a given item will be overlapped by the ERPs elicited by previous and subsequent items, making it difficult to isolate the ERP for each individual stimulus. This is illustrated in Figure 2. We eliminated this overlap problem in the present experiment by using an *irrelevant-probe technique* that was developed previously in ERP studies of spatial attention (Heinze et al., 1990; Luck et al., 1993; Luck




Figure 2. Example of rapid serial visual presentation stream used in Experiment 1. All items were successively presented at fixation. Half the trials had a probe present behind the second target (T2); the other half had no probe present. The event-related potential overlap problem is simulated in the lower portion of the figure.

& Hillyard, 1995). Specifically, a task-irrelevant white square was flashed behind T2 on some trials, and the ERP elicited by this irrelevant probe flash was used as a measure of sensory processing at the time of T2. A subtraction technique was used to separate the ERP elicited by the probe flash from the ERPs elicited by the other items in the RSVP stream, as illustrated in Figure 2. To accomplish this, a probe flash was presented simultaneously with T2 on half of the trials, and T2 was presented without a probe flash on the other half.

The probe technique is based on the fact that voltage fields summate linearly. The ERP waveform recorded on probe-present trials reflects the sum of the response to the probe stimulus and the response to the other items in the RSVP sequence, whereas the ERP waveform recorded on probe-absent trials simply reflects the response to the other items. Consequently, the response to the probe stimulus can be computed by creating a *difference waveform* in which the ERP waveform on probe-absent trials is subtracted from the ERP waveform on probe-present trials.² This technique has been used in several previous studies of spatial attention, which have shown that the probe-elicited P1 and N1 components are suppressed when the probe appears at an ignored location relative to when the probe appears at an attended location (Luck et al., 1993; Luck & Hillyard, 1995).

Method

Participants. All participants in this study were right-handed, neurologically normal college-student volunteers between 18 and

30 years old who were paid or received course credit for their participation. All had normal or corrected-to-normal acuity and reported normal color vision. Twenty individuals participated in Experiment 1.

Stimuli. The stimuli, which are illustrated in Figure 2, were presented on a computer-controlled video monitor placed 70 cm from the participant. Each trial consisted of an RSVP stream of 19 letters and a single digit. Each character in the RSVP stream was presented for 33 ms, and successive characters were separated by a blank interstimulus interval of 50 ms, yielding a presentation rate of approximately 12 characters per second. All characters were displayed individually at the center of a gray screen (6.7 cd/m²). The characters were 0.8° in height and width. The RSVP stream for each trial consisted of 18 nontargets and 2 targets (T1 and T2).

The nontargets in each stream were randomly selected uppercase letters (A–Z, with the exception of Y) drawn in blue (CIE color coordinates: $x = .147, y = .067$). T1 was a digit (2–9), also drawn in blue. Like the nontargets, T2 was an uppercase letter, but it was drawn in red (CIE color coordinates: $x = .636, y = .344$). T1 was equally likely to be an odd or an even digit, and T2 was equally likely to be a vowel or a consonant. On half the trials, a solid white probe square (1° × 1°; 63.2 cd/m²) was presented behind T2, with the same onset and offset time as T2.

Procedure. T1 was either the 7th or the 10th character presented on a given trial. T2 was always the 1st, 3rd, or 7th character presented after T1 (denoted as Lag 1, Lag 3, and Lag 7, respectively). These three lags were used because they typically provide the most relevant data points in the attentional blink function, with little or no impairment at Lags 1 and 7 and the most severe impairment at Lag 3. Five hundred milliseconds after the offset of the final character in the RSVP stream, a question mark appeared for 1,000 ms, signaling the participant to respond. The next trial began 1,000 ms after the offset of the question mark.

Two task conditions were used: single target and dual target. In the dual-target condition, the participants were instructed to make discriminative responses to both T1 and T2 at the end of each trial. Specifically, they made a two-alternative forced-choice response to indicate whether T1 was an even or odd number and a second two-alternative forced-choice response to indicate whether T2 was a vowel or consonant. They were encouraged to make the T1 response before the T2 response, and although either order was allowed, the T1 response was almost always first. In the single-target condition, the participants were instructed to ignore T1 and to make only the vowel–consonant discrimination response at the end of the trial. Each participant received six blocks of 96 trials in each condition; the two conditions were run in counterbalanced order.

Recording and analysis. Electroencephalographic (EEG) activity was recorded from tin electrodes mounted in an elastic cap, located at standard left- and right-hemisphere positions over frontal, central, parietal, occipital, and temporal areas (International 10/20 System sites F3, F4, C3, C4, P3, P4, O1, O2, T3, T4, T5, T6). Two nonstandard sites were used in addition: OL (halfway between O1 and T5) and OR (halfway between O2 and T6). These sites and a right-mastoid site were recorded with a left-mastoid reference, and the data were re-referenced offline to the algebraic average of the left and right mastoids. The horizontal electrooculo-

² The presence of the probe stimulus may alter the processing of the other items in the RSVP stream, and the difference wave created by this subtraction process therefore reflects ERP activity corresponding to this stimulus–stimulus interaction as well as the response to the probe. However, this does not affect the conclusions of this study in any way.

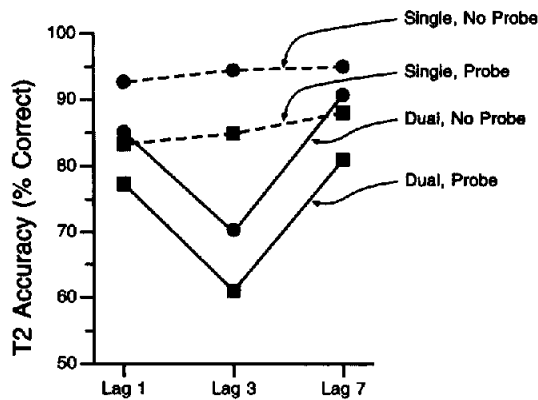


Figure 3. Mean accuracy for identifying the second target (T2) letter in Experiment 1 as a function of lag and probe presence.

gram (EOG) was recorded from electrodes placed 1 cm to the left and right of the external canthi to measure horizontal eye movements, and the vertical EOG was recorded from an electrode beneath the left eye, referenced to the right mastoid, to detect blinks and vertical eye movements. Trials containing ocular artifacts, movement artifacts, or amplifier saturation were excluded from the averaged ERP waveforms. These artifacts led to the rejection of a mean of 9% and a maximum of 21% of trials for a given participant. The EEG and EOG were amplified by an SA Instrumentation amplifier with a bandpass of 0.01–80 Hz (half-power cutoff, Butterworth filters) and were digitized at 250 Hz by a PC-compatible microcomputer.

The amplitude of the P1 and N1 waves was measured as the mean amplitude from 60 to 100 ms and from 140 to 180 ms, respectively, relative to a 200-ms prestimulus baseline. The P1 measurements were obtained at the four most posterior pairs of lateral electrode sites (P3, P4, O1, O2, OL, OR, T5, and T6), and the N1 measurements were obtained at all lateral electrode sites. An analysis of variance (ANOVA) was used for all statistical tests, and all probability values reported were adjusted with the Greenhouse–Geisser epsilon correction for nonsphericity (Jennings & Wood, 1976). To maximize any possible attentional blink effects, for the probe-elicited ERP waveforms shown we included only those trials on which T1 was correctly discriminated. The waveforms include trials with incorrect as well as correct T2 discriminations, however. The behavioral accuracy analyses for T2 were similarly limited to trials on which T1 was correctly discriminated.

Results

Behavior. Figure 3 shows the group mean percentage of trials on which T2 was correctly detected at each lag for the single-target and dual-target conditions. In the dual-target condition, T2 accuracy was approximately equal for Lags 1 and 7 but dropped by 15–20% at Lag 3. This is the typical attentional blink pattern. In the single-target condition, T2 accuracy was approximately equal at Lags 1, 3, and 7, indicating that the impaired accuracy at Lag 3 in the dual-target condition can be attributed to the processing of T1 rather than to some nonspecific factor. In addition, asymptotic T2 accuracy was generally lower in the dual-target condition than in the single-target condition. T2 accuracy was also lower on probe-present trials than on

probe-absent trials, presumably because of masking, but this impairment was similar in magnitude across tasks and stimulus onset asynchronies (SOAs).

To determine the reliability of these effects, we conducted a three-way ANOVA on the T2 accuracy data with the factors of condition (dual target vs. single target), lag (Lag 1, Lag 3, or Lag 7), and probe presence (present vs. absent). This analysis revealed a significant main effect of condition, $F(1, 19) = 113, p < .001$, a significant main effect of lag, $F(2, 38) = 147.32, p < .001$, and a significant Condition \times Lag interaction, $F(2, 38) = 68.54, p < .001$. These effects reflect the generally lower accuracy observed in the dual-target condition and the impairment in accuracy at Lag 3 in the dual-target condition. In addition, there was a significant main effect of probe presence, $F(2, 38) = 27.32, p < .01$, but none of the interactions involving the probe presence factor approached significance. Thus, although the probe flash caused a significant impairment in T2 discrimination accuracy, the probe did not influence the basic attentional blink pattern.

The mean accuracy was 95% correct for T1. There was no significant main effect of lag on T1 accuracy ($F < 1$).

Electrophysiology. The difference waveforms corresponding to the probe flashes are plotted in Figure 4. At the posterior electrode sites, the waveforms contained an initial positive deflection peaking around 90 ms (P1), followed by a negative deflection peaking around 150 ms (N1). The P1 wave was confined to the posterior scalp sites, but the N1 wave could be observed at both anterior and posterior sites. These peaks did not vary significantly in amplitude or latency as a function of lag in either the single- or dual-task




Figure 4. Event-related potential difference waveforms from Experiment 1, calculated by subtracting probe-absent trials from probe-present trials. These waveforms were averaged across participants and across left- and right-hemisphere electrode sites. Note that, by convention, negative is plotted upward. In order to increase visual clarity, all waveforms shown here and in subsequent figures were low-pass-filtered by convolving them with a Gaussian impulse-response function ($SD = 6$ ms; 50% amplitude cutoff at approximately 35 Hz).

condition. Specifically, separate ANOVAs were conducted for the P1 and N1 waves with the factors of condition, T2 lag, and electrode location. No significant main effects or interactions involving the condition or lag factors were obtained for either component. For both components, the only significant effect was a main effect of electrode site: $F(7, 119) = 11.33, p < .005$, for the P1 component, and $F(13, 221) = 36.09, p < .001$, for the N1 component (these significant effects simply reflect the well-established fact that the P1 and N1 components vary in amplitude across scalp sites). It is important that the Condition \times Lag interaction effect that defines the attentional blink did not approach significance for either component ($F_s < 1, p_s > .5$).

Discussion

In this experiment, a normal attentional blink pattern was observed in the behavioral accuracy measures. That is, T2 discrimination accuracy in the dual-task condition was relatively high at Lags 1 and 7, but it decreased significantly at Lag 3. However, despite the substantial impairment in T2 accuracy at Lag 3, there was no suppression of the P1 or N1 components elicited by the probe flash that appeared concurrently with T2. This finding is consistent with the hypothesis that the attentional blink reflects a relatively late suppression of information and stands in sharp contrast to the results of spatial attention experiments, in which the P1 and N1 waves are typically suppressed for stimuli presented at ignored locations (e.g., Eimer, 1994; Luck et al., 1994; Mangun & Hillyard, 1990, 1991).

It is unlikely that the absence of a significant suppression effect in this experiment was caused by a lack of power. First, the effect of the attentional blink on behavioral accuracy in this experiment was at least as large as the attention effects observed in most previous ERP studies of spatial attention, indicating that the attentional manipulation was highly effective. Second, our sample size of 20 participants was larger than the sample sizes used in most previous ERP studies of spatial attention. However, it is always difficult to draw conclusions on the basis of a single null effect. Experiment 2 was therefore conducted to provide converging evidence.

Experiment 2

Although the finding of no P1 or N1 suppression during the attentional blink in Experiment 1 suggests that the attentional blink reflects a postperceptual suppression of information, it is possible that the attentional blink instead reflects suppression at a relatively late substage of perception, after the offset of the P1 and N1 components. Experiment 2 was designed to rule out this possibility by measuring the N400 component elicited by words presented during the attentional blink interval as a means of determining whether these words were identified.

The N400 is a large negative component that peaks approximately 400 ms after the onset of a stimulus and reflects the degree of mismatch between a word and a previously established semantic context (Besson et al.,

1992; Kutas & Hillyard, 1980a; Kutas et al., 1988; Osterhout & Holcomb, 1995; Rugg, 1984). For example, a large N400 would be elicited by the last word in the sentence "The woman drove to work in her shiny new nose" but not by the last word in the sentence "The woman drove to work in her shiny new car." Similar results are obtained for sequentially presented word pairs, in which an N400 is elicited by the second word if it mismatches the semantic context established by the first word (e.g., a large N400 is elicited by the second word in *pickle-robe* but not in *shoe-foot*). Because the meaning of a word must be accessed before it can be compared with a semantic context, the presence of an N400 peak for a mismatching word indicates that the word has been identified to a point at which some elements of the word's meaning are available (although not necessarily all aspects of its meaning). Therefore, the presence of a normal N400 component at the Lag 3 position would provide strong evidence that words presented during the attentional blink are fully identified, even though the observer cannot accurately report them. However, if information is suppressed before perceptual processing is complete during the attentional blink, the N400 should be suppressed for words presented during the blink relative to words presented outside of the blink.

An RSVP paradigm was used for this experiment, with a digit target as T1 and a word target as T2 (see Figure 5A). To isolate the N400 component for the T2 word, we artificially established a semantic context at the beginning of each trial and compared trials on which T2 matched this context with trials on which T2 mismatched this context. Specifically, a context word was presented immediately before each RSVP stream, and the participants were required to report whether the T2 word within the RSVP stream was semantically related or unrelated to this context word. The N400 was then measured from difference waves in which related T2 trials were subtracted from unrelated T2 trials. It is important that the only difference between these two types of trials was the semantic relationship between the context word and the T2 word, and any deviation from zero in the difference waves therefore depended on the accurate identification and semantic analysis of T2. Both the behavioral task and the ERP measurements reflected the same related versus unrelated semantic discrimination. Note that the use of difference waves in this experiment also removed the overlapping sensory ERPs from the previous and subsequent stimuli, which were the same on related T2 trials and unrelated T2 trials.

Method

Participants. Fourteen individuals participated in this experiment. They were drawn from the same population used in Experiment 1 except that they were also required to be native English speakers.

Stimuli. As shown in Figure 5A, each trial began with the presentation of a context word for a duration of 1,000 ms, followed by a blank interval of 1,000 ms. A stream of 20 seven-character strings of letters or numbers was then presented at a rate of 12 items per second. Either the 7th or the 10th string served as T1, and this string consisted of a randomly selected digit (2–9), repeated seven

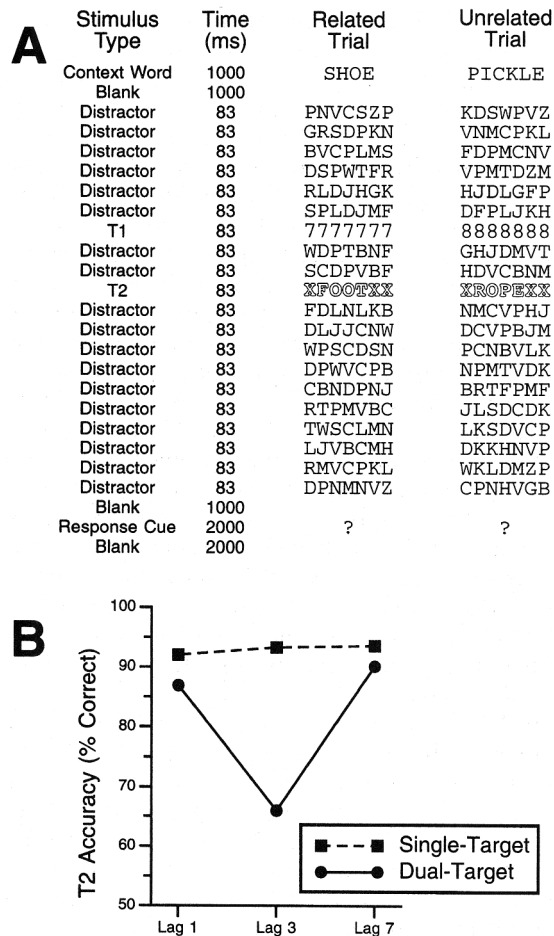


Figure 5. A: Example stimuli from Experiment 2. B: Mean discrimination accuracy for the second target (T2) word as a function of lag for the single- and dual-target conditions in Experiment 2.

times to create a seven-character string. The nontargets in the stream consisted of seven-character strings consisting of randomly selected consonants. All character strings subtended $4.9^\circ \times 0.8^\circ$ of visual angle, and all were presented in blue except for T2, which was presented in red. At the end of each trial, a question mark appeared, signaling the participant to respond. Stimulus timing was identical to that in Experiment 1.

T2 was a word of three to seven characters, drawn in red. Words less than seven characters long were flanked by Xs to create a seven-character string. The T2 word was always the first, third, or seventh string presented after T1 (i.e., Lag 1, Lag 3, or Lag 7). On half the trials, the T2 word was semantically related to the context word (e.g., *doctor-nurse*). On the other half, the T2 word was not semantically related to the context word (e.g., *doctor-chicken*). Each related word pair was randomly selected from a pool of 360 highly related word pairs (Postman & Keppel, 1970), and each unrelated word pair was created by selecting words at random from these pairs.³ Each word was presented twice (once per condition), once within a related word pair and once within an unrelated pair. The ordering of the word pairs was counterbalanced for each participant so that a word used in a related pair in the dual-target condition would be used again in an unrelated pair in the single-target condition. Thus, each T2 word served as both a related

word and an unrelated word and appeared in both conditions of the experiment.

Procedure. In the dual-target condition, the participants were instructed to identify the number string and the red word when they appeared in the RSVP stream and to make 2 two-alternative forced-choice responses at the end of each trial. These responses indicated whether the number string was even or odd and whether the red word was semantically related or unrelated to the context word presented at the beginning of that trial. T1 was equally likely to be an even or an odd number and T2 was equally likely to be related or unrelated to the context word. A single-target condition was also used in this experiment in which the participants were instructed to perform only the T2 word task. Each participant received six 60-trial blocks in the dual-target condition and six 60-trial blocks in the single-target condition (in counterbalanced order).

Recording and analysis. The recording and general analytic procedures were the same as those used in Experiment 1. Artifacts led to the rejection of a mean of 17% and a maximum of 24% of the trials for a given participant.

The averaged ERP waveforms were time locked to the onset of the T2 word. Difference waves were constructed by subtracting the ERP waveforms elicited by T2 words that were related to the context word from the ERP waveforms elicited by T2 words that were unrelated to the context word. The N400 component was measured from these difference waves as the mean amplitude 300–500 ms poststimulus, relative to a 200-ms prestimulus baseline.⁴ Measurements were obtained at the frontal, central, and parietal electrodes (F3, Fz, F4, C3, Cz, C4, P3, Pz, and P4). As in Experiment 1, all behavioral and ERP analyses for T2 were limited to trials on which T1 was correctly discriminated.

Results

Behavior. Mean T2 discrimination accuracy is plotted as a function of T2 lag in Figure 5. Asymptotic accuracy was somewhat higher in the single-target condition than in the dual-target condition. Within the single-target condition, there appeared to be no effect of lag. In the dual-target condition, however, there was a substantial decrease in accuracy for Lag 3 compared with Lags 1 and 7. A two-factor ANOVA was conducted with condition (single vs. dual target) and lag (Lag 1, Lag 3, or Lag 7) as repeated measures factors, and this analysis revealed a significant main effect of condition, $F(1, 13) = 138.51, p < .001$, a significant main effect of lag, $F(2, 26) = 59.00, p < .001$, and a significant Condition \times Lag interaction, $F(2, 26) = 65.71, p < .001$. The mean accuracy for the T1 discrimina-

³ Because unrelated pairs were created by selecting words at random, it is possible that some of the word pairs created in this manner were actually somewhat semantically related. However, such occurrences were rare (if they happened at all) and would have simply decreased overall accuracy and N400 amplitude on unrelated trials. Any such effects would have been small and would have been equivalent across the experimental conditions, leading to no impact on the conclusions drawn from this experiment.

⁴ This measurement window was selected to minimize any potential contributions from the P3 component, which typically follows the N400 component in an experiment of this nature. The N400 was also measured with wider windows, and the results were essentially the same.

tion was 93% correct, with no significant main effect of T2 lag ($F < 1$).

Electrophysiology. The difference waveforms for both the single- and dual-target conditions are plotted as a function of T2 lag in Figure 6. The waveforms consisted primarily of a single large deflection with the usual characteristics of the N400 component (Kutas & Hillyard, 1980b, 1983; Kutas et al., 1988). Specifically, it was a negative-going deflection, peaking at approximately 400 ms, with a centroparietal distribution and slightly greater amplitude over the right hemisphere, especially at the anterior electrode sites.⁵

As shown in Figure 6, the mean N400 amplitude was generally smaller in the dual-target condition than in the single-target condition. However, there was no obvious effect of lag for either condition. A three-way ANOVA on N400 amplitude with the factors of condition, lag, and electrode site yielded a significant main effect of condition, $F(1, 13) = 8.70, p < .01$, corresponding to the generally greater N400 amplitudes observed in the single-target condition. In addition, this difference was greater at sites where the N400 was larger, which led to a significant interaction between lag and electrode site, $F(8, 104) = 8.95, p < .001$. However, the main effect of lag was not significant ($F < 1$), nor were the Lag \times Condition interaction, $F(2, 26) = 1.02, p > .35$, the Lag \times Electrode Site interaction, $F(16, 208) = 1.22, p > .30$, and the three-way interaction among lag,

condition, and electrode site ($F < 1$). Thus, there was no evidence of N400 suppression during the attentional blink period.

Discussion

In Experiment 2, we observed a normal attentional blink pattern for T2 accuracy. That is, T2 accuracy was relatively high at Lags 1 and 7, with a substantial decrease at Lag 3 in the dual-target condition. However, despite this substantial impairment in behavioral performance, there was no suppression of the N400 component at Lag 3. This finding provides strong evidence that the T2 word was identified to the point of meaning extraction⁶ during the attentional blink period because no N400 activity would have been possible in the unrelated-minus-related difference waves unless the meaning of the T2 word had been extracted. The attentional blink therefore appears to reflect a loss of information after stimulus identification is complete, perhaps at the stage of working memory.

Several alternative explanations must be considered before accepting this conclusion. First, it is possible that the difference waves used to measure the N400 component did not actually reflect the N400 component but instead reflected a different component such as the P3 wave or some combination of components. However, our conclusions do not in any way depend on establishing that the activity recorded in this experiment is the same as the N400 activity observed in prior experiments. The simple fact that there was a difference between related T2 trials and unrelated T2 trials during the attentional blink period is sufficient to indicate that the T2 words were identified to the point of meaning extraction. Note, however, that the polarity, timing, and scalp distribution of the activity observed in this experiment are consistent with the hypothesis that this activity consisted primarily of the same N400 component that has been observed previously (Kutas & Hillyard, 1980b; Kutas & Hillyard, 1983; Kutas et al., 1988). We therefore continue to refer to this activity as the N400 component.

A second alternative explanation for our results is that the absence of N400 suppression during the attentional blink reflects a lack of sensitivity or power. Several pieces of evidence argue against this possibility, however. First, the N400 was significantly smaller in the dual-target condition than in the single-target condition, which may reflect a

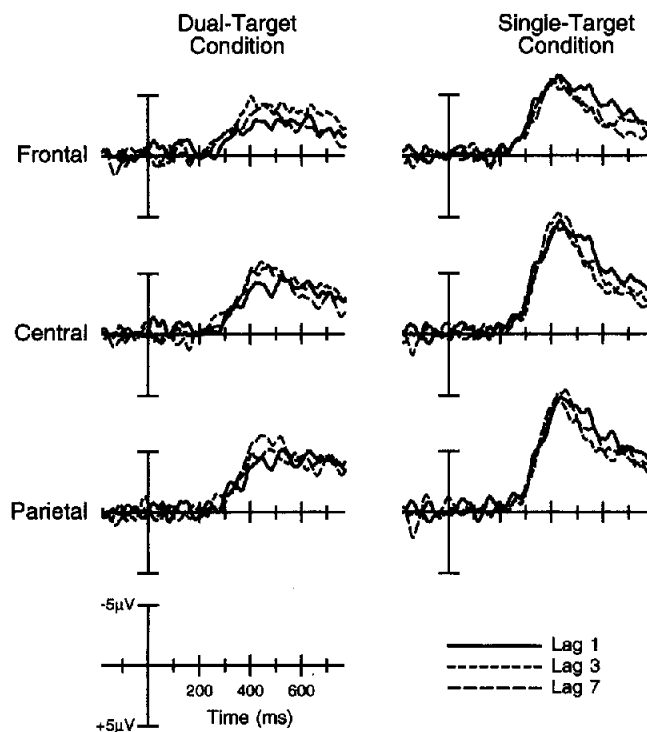


Figure 6. Grand average event-related potential difference waveforms from Experiment 2, formed by subtracting related T2 trials from unrelated T2 trials. These waveforms were recorded at midline electrode sites and were averaged across participants. Negative is plotted upward. T2 = second target.

⁵ The lateralization of the N400 was tested statistically in a four-way ANOVA with the factors of condition (single target or dual target), lag (1, 3, or 7), anterior-posterior electrode position (frontal, central, or parietal), and hemisphere (left or right). Although the main effect of hemisphere did not approach significance, the interaction between anterior-posterior electrode position and hemisphere was significant, $F(2, 26) = 7.22, p < .02$.

⁶ The presence of an N400 does not necessarily imply that a word's meaning has been fully computed. For example, it is possible that the N400 reflects some aspect of lexical access that proceeds more quickly or effortlessly for words that have already been primed. However, the presence of an N400 does indicate that a word has been processed to a level at which meaning-based priming plays a role.

general impairment in processing associated with performing two tasks at the same time. This result indicates that the N400 component is sensitive to at least some cognitive manipulations, and the fact that this difference was significant at the .01 level demonstrates that the signal-to-noise ratio of the ERP waveforms was reasonably high. In addition, previous studies of spatial attention have shown that the N400 is suppressed for words presented at ignored locations (McCarthy & Nobre, 1993; Otten, Rugg, & Doyle, 1993), which also indicates that the N400 is sensitive to some attentional manipulations. However, it is possible that the attentional blink produced a modest amount of perceptual degradation that was sufficient to produce a significant decrement in accuracy but was not sufficient to yield a significant decrease in N400 amplitude. For example, there could be a nonlinear relationship between N400 amplitude and perceptual discriminability such that a constant N400 amplitude is observed unless the perceptual degradation is nearly complete. To our knowledge, there are no data in the literature that quantify the relationship between N400 amplitude and perceptual quality under conditions like those of the present experiment. Therefore, this alternative explanation of our results is difficult to evaluate without additional data. Experiment 3 was designed to provide the necessary data.

Experiment 3

To quantify the relationship between perceptual degradation and N400 amplitude, we used the single-target condition from Experiment 2 and added varying intensities of visual noise to the T2 word stimulus. We predicted that adding visual noise to the word would lead to a substantial decline in both behavioral accuracy and N400 amplitude. By showing that N400 amplitude is reduced when the stimulus is perceptually degraded, we sought to buttress our conclusion that the robust N400 observed during the attentional blink in Experiment 2 can be interpreted as evidence for unimpaired perception.

The single-target condition was used instead of the dual-target condition for this experiment so that we could assess the effects of perceptual degradation without any additional effects of attentional suppression. This necessitated a small change in terminology: Because the "T2 word" was the only target used in this experiment, we simply refer to it as the "target word."

Method

Ten individuals drawn from the same population used in Experiment 2 participated in this experiment. The stimuli, procedure, and recording protocol were identical to those used in the single-target condition of Experiment 1 with the following exceptions. The target word was accompanied by a visual noise array that consisted of 210 dots, each of which measured approximately $0.07^\circ \times 0.07^\circ$. The dots were the same red hue as the target word and were randomly distributed across the same $4.9^\circ \times 0.8^\circ$ rectangle that contained the target word. The amount of perceptual degradation was varied by using three different levels of visual

noise luminance: dim (3.7 cd/m^2), medium (5.3 cd/m^2), and bright (7.2 cd/m^2). Trials without visual noise were also included.

The RSVP streams in this experiment were exactly like those used in Experiment 2. Thus, each trial contained a digit string even though this item was never relevant. In addition, there was a lag of 1, 3, or 7 between the digit string and the target word, but this lag was irrelevant for the present purposes and the data were therefore collapsed across lags.

The different levels of visual noise were equiprobable and were randomly intermixed within trial blocks. Each participant received eight blocks of 60 trials, which yielded the same number of trials per condition as in Experiment 2. As in Experiment 2, the participants were required to report whether the target word was semantically related or unrelated to the context word for that trial. Related and unrelated target words were equiprobable. Artifacts led to the rejection of a mean of 15% and a maximum of 23% of the trials for a given participant.

Results and Discussion

The behavioral and electrophysiological results from Experiment 3 are summarized in Figure 7. As the luminance of the visual noise increased, accuracy in reporting the semantic relationship between the target word and the context word decreased. Note that the accuracy difference between the no-noise and bright-noise trials was comparable to the accuracy difference between Lags 1 and 3 in Experiment 2. The accuracy values were analyzed in an ANOVA with a single factor of noise intensity. This analysis indicated that the effects of visual noise were highly significant, $F(3, 27) = 29.66, p < .001$.

The N400 component observed in this experiment was highly similar to the N400 component observed in Experiment 2 in terms of latency and scalp distribution. However, N400 amplitude was much more closely related to behavioral accuracy in this experiment. Specifically, as the intensity of the visual noise increased, N400 amplitude declined in a manner that closely paralleled the effects of visual noise on accuracy. Statistical support for the effect of visual noise on N400 amplitude was obtained in a two-way ANOVA with the factors of noise intensity and electrode site, which yielded a significant main effect of noise intensity, $F(3, 27) =$




Figure 7. Comparison of N400 amplitude (averaged across scalp sites) and behavioral accuracy as a function of visual noise intensity in Experiment 3. The grand average event-related potential waveforms at the right of the figure were recorded at the midline central electrode site (Cz). Negative is plotted upward.

8.07, $p < .005$. Because N400 amplitude was greatest at the midline central and parietal electrodes, there was also a significant main effect of electrode site, $F(8, 72) = 10.76$, $p < .001$, and a significant interaction between noise intensity and electrode site, $F(24, 216) = 3.67$, $p < .02$. These results demonstrate that the N400 component is highly sensitive to manipulations of perceptual quality. Thus, the lack of a decrement in N400 amplitude during the attentional blink in Experiment 2 can be taken as strong evidence for a postperceptual effect of attention.

Experiment 4

By demonstrating that stimuli are fully identified during the attentional blink, the preceding experiments provided a lower bound on the stage of processing influenced by the attentional blink. In other words, processing is unimpaired at least to the stage of word recognition during the attentional blink. The goal of Experiment 4 was to provide an upper bound by identifying a stage of processing at which processing is impaired during the blink interval. In particular, we examined the P3 component, which has been hypothesized to reflect working memory processes.

The P3 wave is a large positive component that typically peaks 400–600 ms poststimulus and is largest over central and parietal midline sites. The amplitude of the P3 is modulated by the frequency of a target category, with infrequent targets eliciting larger amplitudes than frequent targets (Donchin, 1981). For example, if an observer is required to discriminate between male and female names within a sequence composed of 85% male names and 15% female names, the female names will elicit a much larger P3 component than will the male names. The fact that the P3 component is sensitive to the probability of a task-defined stimulus category implies that it is elicited after the stimulus has been categorized and therefore reflects a postperceptual process. The P3 wave is also sensitive to perceptual manipulations, such as changes in stimulus discriminability (e.g., McCarthy & Donchin, 1983), but this is sensible given that the quality of the input to postperceptual processes will influence those processes.

The most widely accepted theory of the P3 component proposes that it reflects the updating of information in working memory (Donchin, 1981; Donchin & Coles, 1988). There are alternative hypotheses as well (e.g., Verleger, 1988), but virtually all theories of the P3 component would posit that the P3 wave is present only for stimuli that have reached the level of working memory. The presence of an unsuppressed P3 wave during the attentional blink would therefore indicate that the attentional blink occurs after information reaches working memory. In contrast, a suppression of the P3 wave would indicate that the attentional blink occurs at or before the stage of working memory.

In this experiment, we used an RSVP task with individually presented letter and digit stimuli, as in Experiment 1. To isolate the probability-sensitive P3 component, we used a two-alternative forced-choice task in which one alternative occurred on 85% of trials and the other alternative occurred on 15% of trials. The infrequent alternative was expected to

elicit a much larger P3 component than the frequent alternative, and the P3 component was therefore isolated by means of difference waves in which the ERP waveform elicited by the frequent alternative was subtracted from the ERP waveform elicited by the infrequent alternative. This subtraction procedure also eliminated the overlapping sensory responses elicited by the preceding and subsequent stimuli, as in the preceding experiments.

Method

Eight individuals drawn from the same population used in Experiment 1 participated in this experiment. The stimuli, procedure, and recording protocol were identical to those used in Experiment 1 with the following exceptions. All characters appeared in black except for T2, which was white. As in Experiment 1, the distractor items were randomly selected letters, and T1 was the sole digit in the RSVP stream. The participants were required to make a two-alternative forced-choice response at the end of the trial, indicating whether the T1 digit was odd or even (both alternatives were equally likely to occur). The T2 letter was the letter E on 15% of trials and some other letter, selected at random, on 85% of trials. At the end of the trial, the participants were required to press a button if T2 was the infrequent letter E and to make no response to T2 if it was some other letter.⁷ Note that although the letter E served as the infrequent category, this letter was actually more frequent than any individual letter in the frequent category; however, it is the probability of the task-defined category rather than the actual stimulus that determines P3 amplitude (see, e.g., Courchesne, Hillyard, & Courchesne, 1977). Responses to both T1 and T2 were required in the dual-target condition, but only the T2 response was required in the single-target condition. No probe flashes were presented in this experiment.

The averaged ERP waveforms were time locked to the onset of the T2 letter. Difference waves were constructed by subtracting the ERP waveforms elicited by the frequent T2 category from the ERP waveforms elicited by the infrequent T2 category. The P3 component was measured from these difference waves as the mean amplitude 400–800 ms poststimulus, relative to a 200-ms prestimulus baseline. Measurements were obtained at the central and parietal electrodes, where the P3 component was generally largest. A probability-sensitive P2 component was also present in this experiment, and this component was measured as the mean amplitude between 200 and 300 ms at the frontal and central electrode sites. All behavioral and ERP analyses for T2 were limited to trials on which T1 was correctly discriminated. Artifacts led to the rejection of a mean of 14% and a maximum of 26% of the trials for a given participant.

Results

Behavior. T2 accuracy is plotted as a function of lag in Figure 8. As in Experiments 1 and 2, accuracy was slightly lower in the dual-target condition than in the single-target condition at Lags 1 and 7, but it dropped substantially at Lag 3 in the dual-target condition. A two-factor ANOVA was performed with condition (single vs. dual target) and lag

⁷ This go/no-go procedure was used to encourage the participants to use a strategy of making a categorical E/not-E decision rather than simply identifying the letter and performing the task-defined categorization at the end of the trial.

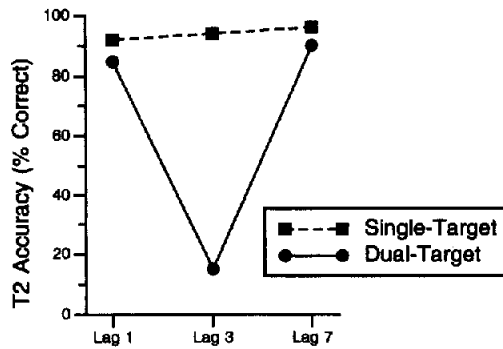


Figure 8. Mean discrimination accuracy for the second target (T2) as a function of lag for the single- and dual-target conditions in Experiment 4.

(Lag 1, 3, or 7) as repeated measures factors. This analysis revealed a significant main effect of condition, $F(1, 7) = 26.04, p < .01$, a significant main effect of lag, $F(2, 14) = 146.04, p < .001$, and a significant Condition \times Lag interaction, $F(2, 14) = 73.63, p < .001$. The mean T1 accuracy was 97% correct, with no significant effect of lag ($F < 1$).

Electrophysiology. The infrequent-minus-frequent difference waveforms are shown in Figure 9. At the central and parietal midline electrode sites, the waveforms primarily

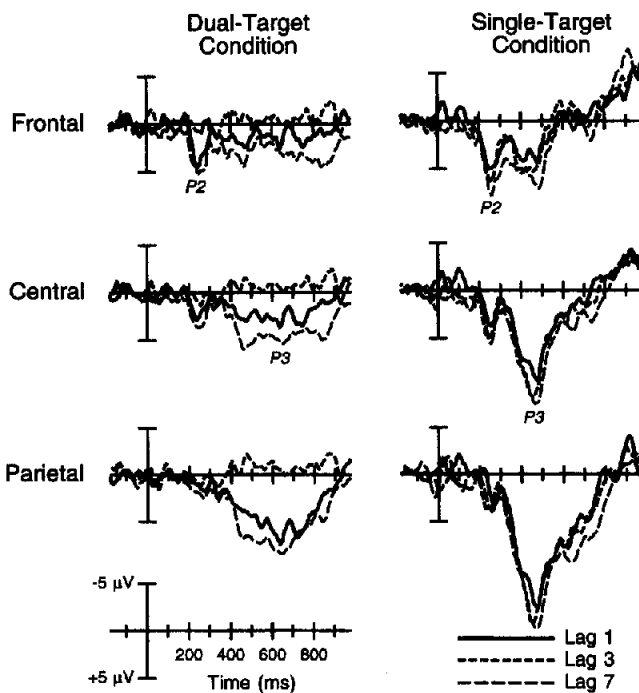


Figure 9. Grand average event-related potential difference waveforms from Experiment 4, formed by subtracting trials with the frequent second target (T2) stimulus from trials with the rare T2 stimulus. These waveforms were recorded at midline electrode sites and were averaged across participants. Negative is plotted upward.

consisted of a large positive deflection peaking around 450 ms (the P3 component). The P3 was generally smaller in the dual-target condition than in the single-target condition. In addition, the P3 wave was unaffected by lag in the single-target condition, but it was completely eliminated at Lag 3 in the dual-target condition, with a mean amplitude of approximately 0 μV . To assess the reliability of these differences, we conducted a three-way ANOVA on P3 amplitude with the factors of condition, T2 lag, and electrode site. This analysis yielded a significant main effect of condition, $F(1, 7) = 35.03, p < .001$, a significant main effect of T2 lag, $F(2, 14) = 25.53, p < .001$, and a significant Condition \times Lag interaction, $F(2, 14) = 5.43, p < .05$. The P3 was broadly distributed with a midline parietal maximum, which led to a significant main effect of electrode site, $F(5, 35) = 23.62, p < .001$, and significant Lag \times Electrode Site, $F(10, 70) = 6.39, p < .01$, and Condition \times Electrode Site, $F(5, 35) = 25.73, p < .001$, interactions. The three-way interaction among lag, condition, and electrode site did not reach significance.

At the frontal and central sites, a P2 component was also present in the difference waves, with a peak latency of approximately 250 ms. Like the P3 component, the P2 component was somewhat smaller for the dual-target condition than for the single-target condition, although the condition main effect did not quite reach significance, $F(1, 7) = 5.06, p < .06$. In addition, the P2 component was unaffected by lag in the single-target condition but was highly suppressed at Lag 3 in the dual-target condition, resulting in a significant main effect of lag, $F(2, 14) = 7.80, p < .02$, and a significant Lag \times Condition interaction, $F(2, 14) = 8.27, p < .01$. P2 amplitude varied across scalp sites, which led to a significant main effect of electrode site, $F(5, 35) = 10.87, p < .005$, and significant Lag \times Electrode Site, $F(10, 70) = 4.38, p < .02$, and Condition \times Electrode Site, $F(5, 35) = 4.13, p < .05$, interactions. The three-way interaction among lag, condition, and electrode site did not reach significance.

Discussion

Unlike the P1, N1, and N400 components, the P3 component was completely suppressed during the attentional blink period. If the P3 component is assumed to reflect some process associated with working memory (Donchin, 1981; Donchin & Coles, 1988), then this finding indicates that the attentional blink operates before or during the process of forming a stable representation of the stimulus in working memory. Even without this assumption, however, our results provide an upper bound on the attentional blink in terms of time. Specifically, the P3 wave was clearly suppressed by 400 ms after the onset of the T2 stimulus at Lag 3, and the P2 wave was suppressed as early as 200 ms. These effects indicate that the attentional blink influenced at least some T2-related processes between 200 and 400 ms poststimulus.

The finding of suppression for the P2 and P3 components but not for the N400 component might appear to be paradoxical, because it seems to suggest that earlier compo-

nents (P2 and P3) can be suppressed without any suppression of later components (N400). However, it is important to realize that the semantic comparison task used to examine the N400 component was substantially more difficult than the letter identification task used to examine the P2 and P3 components, which makes it impossible to compare absolute latencies across experiments. When the P3 and N400 components are examined in comparable tasks, the P3 component typically follows the N400 component, which is consistent with the proposal that the P3 suppression observed in the present experiment reflects an impairment of a process that follows the process reflected by the N400 component. It should also be noted that although we were using the N400 component as a marker for the completion of perceptual processing, this component almost certainly reflects a process that follows perception rather than reflecting the perceptual analysis itself. The presence of an unsuppressed N400 component during the attentional blink does not therefore indicate that the attentional blink begins after the N400 component; this result simply indicates that the processes that necessarily precede the N400 component were not suppressed. In contrast, the finding of P2 and P3 suppression beginning 200–400 ms poststimulus does provide an unambiguous upper bound to the onset of suppression within the context of our stimuli and task.

The P2 component has been studied much less than the P3 component, and it is therefore difficult to make even a tentative interpretation of the P2 suppression effect (except in terms of raw latency). It has previously been shown that the P2 component is sensitive to target probability (Luck & Hillyard, 1994), just as the P3 component is, so the presence of the P2 wave in the infrequent-minus-frequent difference waves was expected. However, it is not clear from previous research whether the P2 component reflects a perceptual process or a postperceptual process (see, e.g., Harter & Aine, 1984; Hillyard & Münte, 1984; Kenemans, Kok, & Smulders, 1993; Luck & Hillyard, 1994). Consequently, the finding of a P2 suppression does not allow us to identify the specific psychologically defined stage at which the attentional blink occurs.

In addition to the suppression observed at Lag 3 in the dual-target condition, the P3 component was also somewhat smaller at Lags 1 and 7 in the dual-target condition than in the single-target condition, just as the N400 component was generally smaller in the dual-target condition in Experiment 2. These lag-independent reductions in P3 and N400 amplitude probably reflect a general impairment in performance under dual-task conditions. Previous dual-task experiments have shown similar decreases in P3 amplitude when cognitive processing resources were withdrawn from the P3-eliciting stimulus by increases in the difficulty of a concurrent task (see, e.g., Isreal, Wickens, Chesney, & Donchin, 1980; Kramer, Wickens, & Donchin, 1983).

General Discussion

The Nature of the Attentional Blink

The results of our experiments provide strong evidence that the attentional blink reflects an impairment that arises

after stimulus identification has been completed, probably at the stage of working memory. The specific findings that support this conclusion are summarized in Figure 10, which shows how the amplitudes of various ERP components were affected by the attentional blink. As can be seen in the figure, the P1, N1, and N400 components were unaffected by the attentional blink, but the P3 component was suppressed during the blink period. Thus, stimulus information appears to be fully identified by the perceptual system during the attentional blink interval, but this information cannot be retrieved at the end of a trial, 1–2 s after the stimulus was presented. In addition, although the attentional blink appears to reflect a postperceptual impairment, this impairment begins relatively soon after the presentation of the target (within 200–400 ms of T2 onset under the conditions of Experiment 4). These results strongly support previous models of the attentional blink that propose a postperceptual locus of suppression (Chun & Potter, 1995; Shapiro & Raymond, 1994; Shapiro, Raymond, & Arnell, 1994). These results are also consistent with the finding that stimuli presented during the attentional blink interval can still prime subsequent stimuli (Maki et al., 1997; Shapiro, Driver, et al., 1997) and the finding that the attentional blink is reduced when T2 is an intrinsically meaningful stimulus, such as the observer's own name (Shapiro, Caldwell, & Sorenson, 1997).

The finding of a suppressed P3 wave during the attentional blink also has important implications for the nature of the attentional blink deficit. In particular, this finding indicates that the impairment in T2 processing begins as early as 200–400 ms after T2 onset, which precludes an impairment arising entirely at a late stage (e.g., response selection). This finding may also help to distinguish between the two leading models of the attentional blink: the two-stage model of Chun and Potter (1995) and the interference model of Shapiro, Raymond, and their colleagues (Isaak, Shapiro, & Martin, in press; Shapiro & Raymond, 1994; Shapiro et al., 1994). The two-stage model proposes that all




Figure 10. Mean amplitudes for the P1 and N1 components from Experiment 1, the N400 component from Experiment 2, and the P3 component from Experiment 4. Only the dual-target condition is shown.

items in the RSVP stream are processed to the point of conceptual representations without attention (Stage 1) and that attention is used to consolidate these representations into a durable and reportable form (Stage 2). The attentional blink is therefore seen as a failure of T2 to receive Stage 2 processing when Stage 2 is still occupied with T1. The interference model similarly proposes that the items in the RSVP stream are fully identified, but this model stresses the role of interference among the identified objects in working memory in explaining the attentional blink. For example, erroneous T2 responses tend not to be random guesses but instead tend to be other items from the RSVP stream that were presented around the same time as T2 (Isaak et al., in press; Maki, Couture, Frigen, & Lien, 1997). The attentional blink is therefore explained in terms of interference produced by the storage of T1 in working memory.

The main difference between these models is that the two-stage model proposes that there is a specific process that cannot be applied to T2 during the attentional blink and that T2 consequently fails to reach working memory, whereas the interference model proposes that T1 and T2 both enter working memory but that T2 is lost because of interference caused by T1. Because the P3 wave is thought to reflect the updating of working memory, the finding of a suppressed P3 wave for T2 during the attentional blink suggests that T2 never reaches working memory; this finding therefore tends to favor the two-stage model. However, the two-stage model does not account as directly for the nonrandom pattern of T2 errors observed during the attentional blink period. We therefore propose a hybrid model that involves both interference and two stages (for a related proposal, see Shapiro, Arnell, & Raymond, 1997). Following Potter (1993), we begin by proposing that all items in the RSVP stream are initially stored in a *conceptual short-term memory* (CSTM) buffer after being fully identified (i.e., after reaching a conceptual level of representation). At this stage, the items are not yet available for report and are prone to decay and to replacement by other incoming stimuli (see Enns & Di Lollo, 1997). We further propose that attention serves to consolidate information stored in the CSTM buffer into a reportable and more durable form, which we call *visual working memory* (VWM; see Baddeley, 1986). Following Duncan and Humphreys (1989) and Shapiro et al. (1994), we propose that the selection of items for transfer from CSTM to VWM is based on the degree of match between each representation in CSTM and a target template. Finally, we propose that when attentional processes are engaged in the process of transferring T1 from CSTM to VWM, they are unavailable for transferring T2, resulting in errors in the report of T2. It is important, however, that these errors during the attentional blink period not be random guesses but reflect the current contents of CSTM. Specifically, as soon as attention becomes available after the consolidation of T1, the system will attempt to transfer any remaining information from CSTM into VWM so that T2 can be reported. Because of decay and interference within CSTM, however, the wrong item will frequently be transferred from CSTM to VWM, resulting in intrusions from other items in the RSVP stream. In this manner, the attentional blink

reflects a problem in retrieving a representation of T2 that had been accurately generated but failed to be stored in a durable and reportable (i.e., retrievable) form.

Multiple Mechanisms of Attention

Our findings contrast strikingly with previous electrophysiological studies of spatial attention, which have generally shown that manipulations of attention can influence processing at an early stage, before stimulus identification is complete. In particular, when attention is directed to one location, the P1, N1, and N400 components are suppressed for stimuli presented at other locations (e.g., Luck et al., 1993, 1994; Mangun & Hillyard, 1988; Mangun et al., 1993; McCarthy & Nobre, 1993; Otten et al., 1993). This difference between time-based and space-based attention effects is consistent with the proposal that space plays a special role in attention (Hillyard & Münte, 1984; Shih & Sperling, 1996; Treisman & Gelade, 1980).

Although space-based attention may begin to influence processing at an earlier stage than other forms of attention, directing attention to a spatial location may involve both perceptual-level attentional mechanisms and the same type of postperceptual attentional mechanisms that were observed in the present study. Consequently, the effects of spatial attention on behavioral output in a given experiment are likely to reflect the combined effects of both early and late selection mechanisms, with the relative contributions of early and late mechanisms varying according to the task. For example, Mangun and Hillyard (1990) examined concurrent changes in behavioral accuracy and ERPs in a paradigm in which the observers varied the allocation of attention between two locations. Their results indicated that behavioral performance was most closely tied to relatively late ERP modulations. Specifically, the P1 and N1 components varied proportionately with the amount of attention directed to the location of the stimulus, indicating that attention could not be allocated to one location without being withdrawn from the other location, whereas the effects of attentional allocation on both the P3 component and behavioral accuracy indicated that the observers could divide attention between the two locations effectively. Thus, behavioral output appeared to correspond to the characteristics of later processing stages rather than earlier processing stages in this specific paradigm. Conversely, recent studies by Suzuki and Cavanagh (1997) and by Cave and Zimmerman (1997) have provided evidence that the detailed structure of behavioral attention effects in other paradigms can be predicted by the attention effects that have been observed at the single-unit level in visual cortex (Luck, Chelazzi, et al., 1997; Moran & Desimone, 1985), indicating that behavioral output was primarily limited by perceptual-level attentional mechanisms in these paradigms.

Our results suggest a model of cognitive processing in which there are two distinct attentional mechanisms, one that is specific to spatial selection and another that is more general. This model is illustrated in Figure 11 in the context of a task with both spatial and temporal components. In this task, spatially distributed five-letter arrays are presented at a

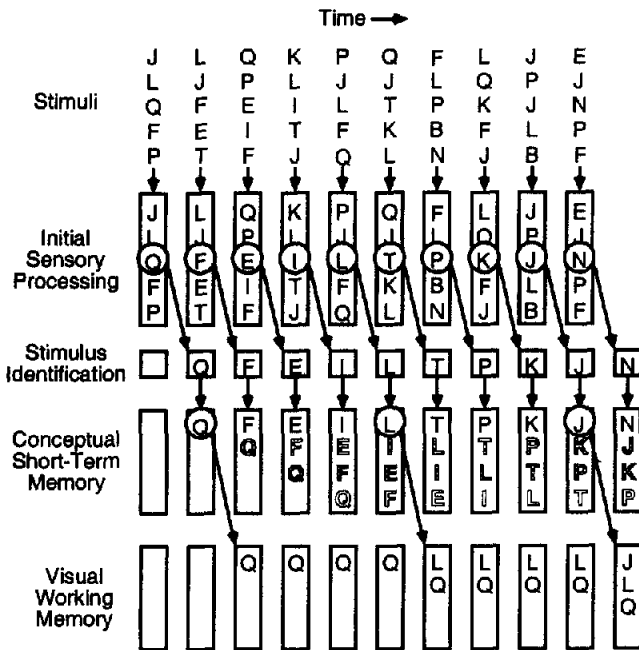


Figure 11. Application of a multistage model of attention to a task with both spatial and temporal components. The observer's task is to identify the central letter in each five-letter stimulus array and to report as many of these letters as possible at the end of the rapid serial visual presentation stream. Spatial attention is used to select the central letter from each array for stimulus identification, which is assumed to occur sufficiently rapidly that each central letter is accurately identified. The visual working memory stage is assumed to be too slow to store each letter that is identified, however, so only every fifth central letter is transferred from conceptual short-term memory into visual working memory.

rapid rate, and the observer must report as many letters as possible from the central location at the end of the trial. Because several sources of evidence indicate that the initial steps in sensory coding occur in parallel and without capacity limitations, the first step in this model consists of an attention-independent sensory decomposition of the basic features in the sensory input. In contrast, the higher level aspects of perceptual processing appear to suffer from interference when multiple objects must be identified in parallel (Luck, Girelli, McDermott, & Ford, 1997; Prinzmetal, Henderson, & Ivry, 1995; Treisman & Schmidt, 1982). This model therefore proposes that spatial attention is used to select a subset of the information from the initial sensory representation for transfer to higher level perceptual processes. This is shown in Figure 11 as the selection of the central letter for transfer from the initial sensory analysis stage to the stimulus identification stage.⁸ Although we propose that the stimulus identification stage is subject to interference from simultaneously presented distractor items, this stage appears to be fast enough to identify simple stimuli such as letters and words at the rates typically used in RSVP tasks (Potter, 1976), at least when attention is used to reduce interference from simultaneously presented distractor items. This is illustrated in Figure 11 as the successful identifica-

tion of the central letter in each stimulus array. As discussed earlier, we also propose that the representations that are computed at this stage remain active for at least several hundred milliseconds in a CSTM buffer. However, these representations are prone to interference and masking and are not available for report until they are transferred into VWM. This process appears to be too sluggish to store each of the perceived letters when the presentation rate is high, and only a small fraction of the information in CSTM can be transferred into VWM. Attention is thus used at this stage to control the transfer process so that only the most relevant information is stored in VWM. Note that it is possible to hold several items in VWM without interference (Luck & Vogel, 1997), but the process of encoding or consolidating an item in working memory is proposed to be subject to severe capacity limitations (see Isaak et al., in press; Luck & Vogel, 1997; Shapiro et al., 1994).

The only novel attribute of this model is that it explicitly combines a perceptual-level attentional mechanism of the type envisioned by early-selection theorists (e.g., Cave & Wolfe, 1990; Treisman, 1996) with a postperceptual attentional mechanism of the type envisioned by late-selection theorists (e.g., Bundesen, 1990; Duncan, 1980; Duncan & Humphreys, 1989). It may seem unparsimonious to postulate two separate attentional mechanisms, but such an architecture may be the only way to account for the many different effects of visual attention that have been described over the past several decades. In addition, given the tremendous complexity of the human brain, it should not be surprising that several different selective mechanisms are necessary. Indeed, there is already evidence for yet a third attentional mechanism that operates at the stage of response selection (see Johnston, McCann, & Remington, 1996; Pashler, 1989).

This model may also help to resolve a recent controversy over the time course of attention in visual search. Specifically, many models of attention propose that attention shifts from item to item approximately every 50 ms during visual search (e.g., Treisman & Gelade, 1980; Wolfe, 1994), but recent studies using a variant of the attentional blink paradigm have indicated that shifts of attention require approximately 500 ms, a difference of an order of magnitude (Duncan et al., 1994; Ward, Duncan, & Shapiro, 1996). This discrepancy could be resolved by postulating that the perceptual-level spatial attention mechanism can shift quickly from item to item during visual search, whereas the postperceptual attentional mechanism that operates in the attentional blink paradigm is substantially more sluggish (for additional discussion of this issue, see Moore, Egeth, Berglan, & Luck, 1996). Additional research is necessary to determine whether this explanation is correct.

⁸ Note that this attentional mechanism is shown as selecting a single object for identification in Figure 11, but it is entirely possible that this stage operates in a limited-capacity parallel manner rather than a strictly serial manner. In addition, there may be conditions under which multiple items can be identified in parallel without any capacity limitations (see Cohen & Ivry, 1991; Cohen & Rafal, 1991; Luck, Girelli, et al., 1997).

Perception Without Awareness?

Our results also speak to the general issue of perception without awareness (Greenwald, Draine, & Abrams, 1996; He, Cavanagh, & Intriligator, 1996; Jacoby & Whitehouse, 1989; Kunst-Wilson & Zajonc, 1980). The finding of no reduction of the P1, N1, and N400 components despite a substantial impairment in behavioral accuracy during the attentional blink suggests that it is possible for the brain to process a stimulus to the point of lexical access and meaning extraction without storing it in a form that can be accessed for overt report. This conclusion is especially strong in the case of the N400 results because the N400 component was measured in difference waves that reflected the same unrelated-versus-related semantic discrimination that was the basis for the observers' overt reports. In addition, the results of Experiment 3 indicate that the absence of an N400 suppression during the attentional blink cannot be explained by an insensitivity of the N400 component to perceptual degradation. Indeed, the amplitude of the N400 component appears to parallel behavioral accuracy closely when the stimuli are degraded perceptually. Thus, these data indicate that it is possible to impair the accuracy of overt report without impairing perceptual processing.

An alternative explanation that must be considered is that the observers were briefly aware of the T2 stimulus during the attentional blink interval but that this information faded by the time of the overt report, 1–2 s later. However, the complete suppression of the P3 component during the attentional blink argues against this possibility. Many previous researchers have examined the P3 component in tasks that use explicit measures of discrimination accuracy, and a robust P3 wave is virtually always observed for low-probability targets that are explicitly detected (for reviews, see Johnson, 1986; Pritchard, 1981). For example, when an observer performs a difficult signal-detection task, a P3 component is observed for both hits and false alarms as long as the observer is confident about the response (Hillyard, Squires, Bauer, & Lindsay, 1971), which indicates that the P3 is more closely related to the observer's ultimate decision about the target than to the actual presence or absence of the target. The complete suppression of the P3 component during the attentional blink therefore suggests that the T2 stimulus failed to reach awareness during the blink interval. It should be noted, however, that behavioral performance did not fall completely to chance levels during the attentional blink. Therefore, our results are not sufficient to demonstrate that perception can occur in the complete absence of awareness, although they do appear to indicate that it is possible to manipulate awareness while leaving perception intact.

Several previous studies have provided evidence for the possibility of perception without awareness (e.g., Greenwald et al., 1996; Jacoby & Whitehouse, 1989; Kunst-Wilson & Zajonc, 1980), and the main contribution of the present study with respect to this issue is that it suggests a reason *why* perception may sometimes occur in the absence of awareness. When combined with previous research using the RSVP paradigm (especially Potter, 1976), our findings

indicate that the visual system may be able to identify stimuli faster than they can be processed by postperceptual systems. This difference in speed makes it reasonable to suppose that perceptual and postperceptual processes can operate somewhat independently. In other words, given the relative sluggishness of the postperceptual systems, visual tasks might be performed more efficiently if the perceptual system could transfer only the most relevant objects that are identified to the postperceptual processes. This idea is conceptually similar to the proposal that preattentive information may be used to guide perceptual-level attentional mechanisms to items that have relevant features (Treisman & Sato, 1990; Wolfe, Cave, & Franzel, 1989), except that we are proposing that post-perceptual-level attentional mechanisms might be guided to task-relevant stimuli that have already been fully identified. Of course, attentional guidance may well occur in both of these ways.

References

- Allport, D. A., Tipper, S. P., & Chmiel, N. R. J. (1985). Perceptual integration and postcategorical filtering. In M. I. Posner & O. S. Marin (Eds.), *Attention and performance XI* (pp. 107–132). Hillsdale, NJ: Erlbaum.
- Baddeley, A. D. (1986). *Working memory*. Oxford, England: Clarendon Press.
- Besson, M., Kutas, M., & Van Petten, C. (1992). An event-related potential (ERP) analysis of semantic congruity and repetition effects. *Journal of Cognitive Neuroscience*, 4, 132–149.
- 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. (1990). A theory of visual attention. *Psychological Review*, 97, 523–547.
- Cave, K. R., & Wolfe, J. M. (1990). Modeling the role of parallel processing in visual search. *Cognitive Psychology*, 22, 225–271.
- Cave, K. R., & Zimmerman, J. M. (1997). Flexibility in spatial attention before and after practice. *Psychological Science*, 8(5), 399–403.
- Chelazzi, L., Miller, E. K., Duncan, J., & Desimone, R. (1993). A neural basis for visual search in inferior temporal cortex. *Nature*, 363, 345–347.
- Chun, M. M., & Potter, M. C. (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.
- Clark, V. P., & Hillyard, S. A. (1996). Spatial selective attention affects early extrastriate but not striate components of the visual evoked potential. *Journal of Cognitive Neuroscience*, 8, 387–402.
- Cohen, A., & Ivry, R. B. (1991). Density effects in conjunction search: Evidence for a coarse location mechanism of feature integration. *Journal of Experimental Psychology: Human Perception and Performance*, 17, 891–901.
- Cohen, A., & Rafal, R. D. (1991). Attention and feature integration: Illusory conjunctions in a patient with a parietal lobe lesion. *Psychological Science*, 2, 106–110.
- Courchesne, E., Hillyard, S. A., & Courchesne, R. Y. (1977). P3 waves to the discrimination of targets in homogenous and heterogeneous stimulus sequences. *Psychophysiology*, 14, 590–597.
- Donchin, E. (1981). Surprise! . . . Surprise? *Psychophysiology*, 18, 493–513.

- Donchin, E., & Coles, M. G. H. (1988). Is the P300 component a manifestation of context updating? *Behavioral and Brain Sciences*, *11*, 357-374.
- Donchin, E., Karis, D., Bashore, T. R., Coles, M. G. H., & Gratton, G. (1986). Cognitive psychophysiology and human information processing. In M. G. H. Coles, E. Donchin, & S. W. Porges (Eds.), *Psychophysiology: Systems, processes and applications* (pp. 244-267). New York: Guilford Press.
- Driver, J., & Tipper, S. P. (1989). On the nonselectivity of "selective" seeing: Contrasts between interference and priming in selective attention. *Journal of Experimental Psychology: Human Perception and Performance*, *15*, 304-314.
- Duncan, J. (1980). The locus of interference in the perception of simultaneous stimuli. *Psychological Review*, *87*, 272-300.
- Duncan, J., & Humphreys, G. (1989). Visual search and stimulus similarity. *Psychological Review*, *96*, 433-458.
- Duncan, J., Ward, R., & Shapiro, K. (1994). Direct measurement of attentional dwell time in human vision. *Nature*, *369*, 313-315.
- Eimer, M. (1994). "Sensory gating" as a mechanism for visuospatial orienting: Electrophysiological evidence from trial-by-trial cuing experiments. *Perception & Psychophysics*, *55*, 667-675.
- Enns, J. T., & Di Lollo, V. (1997). Object substitution: A new form of masking in unattended visual locations. *Psychological Science*, *8*, 135-139.
- Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in nonssearch task. *Perception & Psychophysics*, *16*, 143-149.
- Eriksen, C. W. (1995). The flankers task and response competition: A useful tool for investigating a variety of cognitive problems. *Visual Cognition*, *2*, 101-118.
- Francolini, C. M., & Egeth, H. E. (1980). On the nonautomaticity of "automatic" activation: Evidence of selective seeing. *Perception & Psychophysics*, *27*, 331-342.
- Greenwald, A. G., Draine, S. C., & Abrams, R. L. (1996). Three cognitive markers of unconscious semantic activation. *Science*, *273*, 1699-1702.
- Hagenaar, R., & van der Heijden, A. H. C. (1986). Target-noise separation in visual selective attention. *Acta Psychologica*, *62*, 161-176.
- Harter, M. R., & Aine, C. J. (1984). Brain mechanisms of visual selective attention. In R. Parasuraman & D. R. Davies (Eds.), *Varieties of attention* (pp. 293-321). London: Academic Press.
- He, S., Cavanagh, P., & Intriligator, J. (1996). Attentional resolution and the locus of visual awareness. *Nature*, *383*, 334-337.
- Heinze, H. J., Luck, S. J., Mangun, G. R., & Hillyard, S. A. (1990). Visual event-related potentials index focused attention within bilateral stimulus arrays: I. Evidence for early selection. *Electroencephalography and Clinical Neurophysiology*, *75*, 511-527.
- Heinze, H. J., Mangun, G. R., Burchert, W., Hinrichs, H., Scholz, M., Münte, T. F., Gös, A., Scherg, M., Johannes, S., Hundeshagen, H., Gazzaniga, M. S., & Hillyard, S. A. (1994). Combined spatial and temporal imaging of brain activity during visual selective attention in humans. *Nature*, *372*, 543-546.
- Hillyard, S. A., Anllo-Vento, L., Clark, V. P., Heinze, H. J., Luck, S. J., & Mangun, G. R. (1996). Neuroimaging approaches to the study of visual attention: A tutorial. In A. F. Kramer, M. G. H. Coles, & G. D. Logan (Eds.), *Converging operations in the study of visual selective attention* (pp. 107-138). Washington, DC: American Psychological Association.
- Hillyard, S. A., Mangun, G. R., Woldorff, M. G., & Luck, S. J. (1995). Neural systems mediating selective attention. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences* (pp. 665-681). Cambridge, MA: MIT Press.
- Hillyard, S. A., & Münte, T. F. (1984). Selective attention to color and location: An analysis with event-related brain potentials. *Perception & Psychophysics*, *36*, 185-198.
- Hillyard, S. A., & Picton, T. W. (1987). Electrophysiology of cognition. In F. Plum (Ed.), *Handbook of physiology: Section 1. The nervous system: Vol. 5. Higher functions of the brain, Pt. 2* (pp. 519-584). Bethesda, MD: Waverly Press.
- Hillyard, S. A., Squires, K. C., Bauer, J. W., & Lindsay, P. H. (1971). Evoked potential correlates of auditory signal detection. *Science*, *172*, 1357-1360.
- Isaak, M. I., Shapiro, K. L., & Martin, M. J. (in press). The attentional blink reflects retrieval competition among multiple RSVP items: Tests of the interference model. *Journal of Experimental Psychology: Human Perception and Performance*.
- Isreal, J. B., Wickens, C. D., Chesney, G. L., & Donchin, E. (1980). The event-related brain potential as an index of display-monitoring workload. *Human Factors*, *22*, 211-224.
- Jacoby, L. L., & Whitehouse, K. (1989). An illusion of memory: False recognition influenced by unconscious perception. *Journal of Experimental Psychology: General*, *118*, 126-135.
- Jennings, J. R., & Wood, C. C. (1976). The e-adjustment procedure for repeated-measures analyses of variance. *Psychophysiology*, *13*, 277-278.
- Johnson, R., Jr. (1986). A triarchic model of P300 amplitude. *Psychophysiology*, *23*, 367-384.
- Johnston, J. C., McCann, R. S., & Remington, R. W. (1996). Selective attention operates at two processing loci. In A. Kramer & G. Logan (Eds.), *Converging operations in the study of visual selective attention* (pp. 439-458). Washington, DC: American Psychological Association.
- Kenemans, J. L., Kok, A., & Smulders, F. T. Y. (1993). Event-related potentials to conjunctions of spatial frequency and orientation as a function of stimulus parameters and response requirements. *Electroencephalography and Clinical Neurophysiology*, *88*, 51-63.
- Kim, M.-S., & Cave, K. R. (1995). Spatial attention in visual search for features and feature conjunctions. *Psychological Science*, *6*, 376-380.
- Kramer, A. F., Wickens, C. D., & Donchin, E. (1983). An analysis of the processing requirements of a complex perceptual-motor task. *Human Factors*, *25*, 597-621.
- Kunst-Wilson, W. R., & Zajonc, R. B. (1980). Affective discrimination of stimuli that cannot be recognized. *Science*, *207*, 557-558.
- Kutas, M., & Hillyard, S. A. (1980a). Event-related brain potentials to semantically inappropriate and surprisingly large words. *Biological Psychology*, *11*, 99-116.
- Kutas, M., & Hillyard, S. A. (1980b). Reading senseless sentences: Brain potentials reflect semantic incongruity. *Science*, *207*, 203-205.
- Kutas, M., & Hillyard, S. A. (1983). Event-related brain potentials to grammatical errors and semantic anomalies. *Memory & Cognition*, *11*, 539-550.
- Kutas, M., Van Petten, C., & Besson, M. (1988). Event-related potential asymmetries during the reading of sentences. *Electroencephalography and Clinical Neurophysiology*, *69*, 218-233.
- Lavie, N. (1995). Perceptual load as a necessary condition for selective attention. *Journal of Experimental Psychology: Human Perception and Performance*, *21*, 451-468.
- Luck, S. J. (1998). Neurophysiology of selective attention. In H. Pashler (Ed.), *Attention* (pp. 257-295). London: UCL Press.
- Luck, S. J., Chelazzi, L., Hillyard, S. A., & Desimone, R. (1997). Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *Journal of Neurophysiology*, *77*, 24-42.
- Luck, S. J., Fan, S., & Hillyard, S. A. (1993). Attention-related

- modulation of sensory-evoked brain activity in a visual search task. *Journal of Cognitive Neuroscience*, 5, 188–195.
- Luck, S. J., Girelli, M., McDermott, M. T., & Ford, M. A. (1997). Bridging the gap between monkey neurophysiology and human perception: An ambiguity resolution theory of visual selective attention. *Cognitive Psychology*, 33, 64–87.
- Luck, S. J., & Hillyard, S. A. (1994). Electrophysiological correlates of feature analysis during visual search. *Psychophysiology*, 31, 291–308.
- Luck, S. J., & Hillyard, S. A. (1995). The role of attention in feature detection and conjunction discrimination: An electrophysiological analysis. *International Journal of Neuroscience*, 80, 281–297.
- Luck, S. J., Hillyard, S. A., Mouloua, M., Woldorff, M. G., Clark, V. P., & Hawkins, H. L. (1994). Effects of spatial cuing on luminance detectability: Psychophysical and electrophysiological evidence for early selection. *Journal of Experimental Psychology: Human Perception and Performance*, 20, 887–904.
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390, 279–281.
- 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.
- Mangun, G. R., & Hillyard, S. A. (1988). Spatial gradients of visual attention: Behavioral and electrophysiological evidence. *Electroencephalography and Clinical Neurophysiology*, 70, 417–428.
- Mangun, G. R., & Hillyard, S. A. (1990). Allocation of visual attention to spatial location: Event-related brain potentials and detection performance. *Perception & Psychophysics*, 47, 532–550.
- Mangun, G. R., & Hillyard, S. A. (1991). Modulations of sensory-evoked brain potentials indicate changes in perceptual processing during visual-spatial priming. *Journal of Experimental Psychology: Human Perception and Performance*, 17, 1057–1074.
- Mangun, G. R., Hillyard, S. A., & Luck, S. J. (1993). Electrocortical substrates of visual selective attention. In D. Meyer & S. Kornblum (Eds.), *Attention and performance XIV* (pp. 219–243). Cambridge, MA: MIT Press.
- McCarthy, G., & Donchin, E. (1983). Chronometric analysis of human information processing. In A. W. K. Gaillard & W. Ritter (Eds.), *Tutorials in ERP research: Endogenous components* (pp. 251–268). Amsterdam: North-Holland.
- McCarthy, G., & Nobre, A. C. (1993). Modulation of semantic processing by spatial selective attention. *Electroencephalography and Clinical Neurophysiology*, 88, 210–219.
- Moore, C. M., Egeth, H., Berglan, L. R., & Luck, S. J. (1996). Are attentional dwell times inconsistent with serial visual search? *Psychonomic Bulletin & Review*, 3, 360–365.
- Moran, J., & Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science*, 229, 782–784.
- Neville, H. J., & Lawson, D. (1987). Attention to central and peripheral visual space in a movement detection task: I. Normal hearing adults. *Brain Research*, 405, 253–267.
- Nissen, M. J. (1985). Accessing features and objects: Is location special? In M. I. Posner & O. S. M. Marin (Eds.), *Attention and performance XI* (pp. 205–219). Hillsdale, NJ: Erlbaum.
- Osterhout, L., & Holcomb, P. J. (1995). Event-related potentials and language comprehension. In M. D. Rugg & M. G. H. Coles (Eds.), *Electrophysiology of mind* (pp. 171–215). New York: Oxford University Press.
- Otten, L. J., Rugg, M. D., & Doyle, M. C. (1993). Modulation of event-related potentials by word repetition: The role of visual selective attention. *Psychophysiology*, 30, 559–571.
- Pashler, H. (1989). Dissociations and dependencies between speed and accuracy: Evidence for a two-component theory of divided attention in simple tasks. *Cognitive Psychology*, 21, 469–514.
- Postman, L. J., & Keppel, G. (1970). *Norms of word association*. New York: Academic Press.
- Potter, M. C. (1976). Short-term conceptual memory for pictures. *Journal of Experimental Psychology: Human Learning and Memory*, 2, 509–522.
- Potter, M. C. (1993). Very short-term conceptual memory. *Memory & Cognition*, 21, 156–161.
- Prinzmetal, W., Henderson, D., & Ivry, R. B. (1995). Loosening the constraints on illusory conjunctions: The role of exposure duration and attention. *Journal of Experimental Psychology: Human Perception and Performance*, 21, 1362–1375.
- Pritchard, W. S. (1981). Psychophysiology of P300. *Psychological Bulletin*, 89, 506–540.
- 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.
- Reeves, A., & Sperling, G. (1986). Attention gating in short-term visual memory. *Psychological Review*, 93, 180–206.
- Rugg, M. D. (1984). Event-related potentials and the phonological processing of words and nonwords. *Neuropsychologia*, 22, 435–443.
- Shapiro, K. L., Arnell, K. M., & Raymond, J. E. (1997). The attentional blink. *Trends in Cognitive Science*, 1, 291–295.
- Shapiro, K. L., Caldwell, J. I., & 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., 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 rapid serial visual presentation. *Journal of Experimental Psychology: Human Perception and Performance*, 20, 357–371.
- Shih, S., & Sperling, G. (1996). Is there feature-based attentional selection in visual search? *Journal of Experimental Psychology: Human Perception and Performance*, 22, 758–779.
- 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.
- Treisman, A. M. (1964). Selective attention in man. *British Medical Bulletin*, 20, 12–16.
- Treisman, A. (1996). The binding problem. *Current Opinion in Neurobiology*, 6, 171–178.
- Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12, 97–136.
- Treisman, A., & Sato, S. (1990). Conjunction search revisited. *Journal of Experimental Psychology: Human Perception and Performance*, 16, 459–478.

- Treisman, A., & Schmidt, H. (1982). Illusory conjunctions in the perception of objects. *Cognitive Psychology, 14*, 107-141.
- Tsal, Y., & Lavie, N. (1988). Attending to color and shape: The special role of location in selective visual processing. *Perception & Psychophysics, 44*, 15-21.
- Tsal, Y., & Lavie, N. (1993). Location dominance in attending to color and shape. *Journal of Experimental Psychology: Human Perception and Performance, 19*, 131-139.
- Verleger, R. (1988). Event-related potentials and cognition: A critique of the context updating hypothesis and an alternative interpretation of P3. *Behavioral and Brain Sciences, 11*, 343-427.
- Ward, R., Duncan, J., & Shapiro, K. (1996). The slow time-course of visual attention. *Cognitive Psychology, 30*, 79-109.
- Weichselgartner, E., & Sperling, G. (1987). Dynamics of automatic and controlled visual attention. *Science, 238*, 778-780.
- Wolfe, J. M. (1994). Guided Search 2.0: A revised model of visual search. *Psychonomic Bulletin & Review, 1*, 202-238.
- Wolfe, J. M., Cave, K. R., & Franzel, S. L. (1989). Guided Search: An alternative to the feature integration model for visual search. *Journal of Experimental Psychology: Human Perception and Performance, 15*, 419-433.
- Yantis, S., & Johnston, J. C. (1990). On the locus of visual selection: Evidence from focused attention tasks. *Journal of Experimental Psychology: Human Perception and Performance, 16*, 135-149.

Received April 23, 1997

Revision received October 3, 1997

Accepted October 3, 1997 ■