

Visual masking plays two roles in the attentional blink

JAMIE C. BREHAUT, JAMES T. ENNS, and VINCENT DI LOLLO
University of British Columbia, Vancouver, British Columbia, Canada

When two targets are displayed in rapid visual sequence and masked by trailing patterns, identification accuracy is nearly perfect for the first target but follows a U-shaped pattern over temporal lag for the second target. Three experiments examined the role of visual masking in this *attentional blink*. Experiment 1 compared integration and interruption masks for both targets. Although either mask was effective in producing the blink when applied to the first target, only the interruption mask was effective when applied to the second target. Experiment 2 showed that integration masking of the second target was ineffective over a wide range of accuracy levels. Combining the two forms of masking in Experiment 3 confirmed the dissociation: A combined mask had only a main effect on accuracy for the first target, whereas it produced a qualitatively different pattern over temporal lag for the second target. These results suggest that representations of the target are substituted in consciousness by that of the interruption mask when visual attention is preoccupied.

There is growing interest among vision researchers in a new way of studying the deployment of attention over time (Duncan, Ward, & Shapiro, 1994; Hussain, Shapiro, Martin, & Kennard, 1997; Luck, Vogel, & Shapiro, 1996; Shapiro, 1994). The method and its main result, referred to in shorthand as the *attentional blink* (or AB), can be likened to a psychophysical cyclotron. Stimuli are smashed into the visual system so rapidly that the system breaks down. By looking at the fragments, it is possible to piece together the otherwise impenetrable sequence of operations lying between the processes of early visual registration and the conscious products of perception.

In one popular version of the method, known as *rapid serial visual presentation* (RSVP), observers attempt to identify two target letters inserted into a temporal stream of digits (Chun & Potter, 1995; Raymond, Shapiro, & Arnell, 1992; Shapiro, 1994). All stimuli are displayed in the same location, one every 100 msec or so, such that each target is masked by the next item in the stream. In a second, elegantly simplified method, two targets, masked by trailing pattern masks, are displayed in different screen locations at various temporal lags from each other (Duncan et al., 1994; Ward, Duncan, & Shapiro, 1997). In both cases, detection or identification of the first target is nearly perfect, whereas it is dramatically reduced for the second target. This impairment is greatest when the second target lags the first by about 200–500 msec, such that the accuracy function is characteristically U-shaped with lag.

This second-target deficit in accuracy is the result denoted as the AB.

Several important results have pointed to cognitive mechanisms underlying the AB. For example, both the temporal extent and the magnitude of the AB are increased along with the difficulty of detecting and reporting the first target (Chun & Potter, 1995; Seiffert & Di Lollo, 1997; but see Shapiro, Raymond, & Arnell, 1994, for an opposing view). This clearly implicates limited-capacity processes associated with attending to the first target. Second, the AB becomes even more pronounced when the first and second targets share an abstract identity (e.g., same letter name). This points more specifically to a high-level competition of episodic object *tokens* for the same conceptual semantic *type* (Shapiro, Caldwell, & Sorensen, 1997; Ward et al., 1997). Third, there is behavioral and electrophysiological evidence that the second target in the AB task, which cannot be reported by observers, nonetheless has an influence that is not available to consciousness (Luck et al., 1996; Shapiro, Driver, Ward, & Sorensen, 1997). This also suggests that the interference or competition occurs at fairly high levels of visual processing.

It is also known that visual masking is essential to the AB, possibly in more than one way. This is potentially of great importance, since studies of masking have historically been helpful in shedding light on questions of attention and the time course of visual processing. Masking has also played a central role in understanding visual processes that occur outside of conscious awareness. For these reasons, we undertook a systematic examination of the role of masking in the AB. As context for our experiments, we will first summarize what is known about the relations between masking and attention generally, as well as what is known about the role of masking in the AB more specifically.

This research was supported by grants from the Natural Sciences and Engineering Research Council of Canada to J. T. Enns and V. Di Lollo. Correspondence should be addressed to J. T. Enns, Department of Psychology, University of British Columbia, 2136 West Mall, Vancouver, BC, V6T 1Z4 Canada (e-mail: jenns@psych.ubc.ca).

—Accepted by previous editor, Myron L. Braunstein

Visual Masking and Attention

Pattern masks present the visual system with at least two different kinds of spatiotemporal conflict (Breitmeyer, 1984; Ganz, 1975; Kahneman, 1968; Scheerer, 1973; Turvey, 1973). One conflict occurs when the target and mask patterns are perceived as part of the same unitary pattern. In this case, masking is akin to the addition of spatial noise (the mask) to the signal (the target) at early levels of visual representation. This is often referred to as masking by *integration*. The temporal characteristic of this form of masking is approximate symmetry around a maximum level of masking at a stimulus onset asynchrony (SOA) of 0 msec, with a release from masking beyond an SOA of about 100 msec in each direction.

A second kind of conflict can arise when processing of a first pattern (the target) is interrupted by a second pattern (the mask) that appears in the same spatial location before the target pattern has been fully processed. This conflict does not involve the early visual stages of processing where contours are defined, but instead it involves a competition for the higher level mechanisms involved in object recognition. It is therefore often referred to as masking by *interruption*. The temporal characteristics are also very different from masking by integration. Interruption masking can occur only when the mask follows the target in time. Therefore, the masking function is often referred to as U- or J-shaped, because target accuracy is lowest at SOAs that are greater than zero, before it begins to improve at even longer SOAs (e.g., Bachmann & Allik, 1976; Turvey, 1973).

A finding that is of direct relevance to the AB is that masking by integration and masking by interruption are differentially sensitive to manipulations of spatial attention. For example, presenting the observer with a varying number of potential target items from trial to trial (i.e., a manipulation of set size) has very little effect on measurements of masking by integration. In contrast, increases in set size dramatically increase both the magnitude and the temporal extent of masking by interruption (Breitmeyer, 1984; Scheerer, 1973; Spencer & Shuntich, 1970). Conversely, manipulations of mask and target intensity have strong effects on masking by integration, where target identification is influenced by the overall extent of camouflage. However, these same manipulations have very little influence on masking by interruption (Breitmeyer, 1984; Spencer & Shuntich, 1970). Indeed, in a recent study (Enns & Di Lollo, 1997), it was shown that four small dots that shared no spatial overlap with the target were very effective in producing masking by interruption, provided that spatial attention was not focused on the correct spatial location.

Masking and the Attentional Blink

A review of the AB literature suggests that the main reason for the use of visual masks to date has been methodological convenience. Simply put, masking the first and second targets brings accuracy performance into a range

where the AB can be observed. Without a mask item following each target item, accuracy is close to the ceiling for both targets (Moore, Egeth, Berglan, & Luck, 1996). Yet, rather than being concerned with the relations between targets and the masking items that follow them, theoretical interest has focused on interference in processing that may occur between the two targets (Chun & Potter, 1995; Duncan et al., 1994; Shapiro et al., 1994). In this paper, we demonstrate that the relations between targets and mask are of even greater importance.

A first hint that visual masking plays an important role can be seen in the seminal paper by Raymond et al. (1992), where it was reported that the AB disappeared when the item immediately following the first target in the stream was replaced by a blank interval. The authors' explanation of this effect appealed in general terms to the increased difficulty of identification when a target was immediately followed by another item. This effect has since been pursued in greater detail. Seiffert and Di Lollo (1997) compared three stimulus conditions for the first target: a blank interval following the target, a digit following the target (interruption masking), and a digit superimposed on the target (integration masking). An AB was observed only in the two masking conditions and was similar in magnitude and extent in both. Grandison, Ghirardelli, and Egeth (1997) showed further that the item following the first target had to consist of a pattern, rather than merely a diffuse light transient, in order to produce the AB. It therefore appears that pattern masking of the first target is at least one of the necessary conditions for the AB. Whether that pattern is presented simultaneously (as in integration masking) or subsequently (as in interruption masking) does not matter. Masking the first target with another pattern seems to ensure that sufficient attentional resources are occupied in the processing of the first target so that second-target processing is impaired.

A second hint that visual masking is critical is well known to AB researchers, but not formally documented in the literature. Specifically, in order to obtain an AB, the second target must be followed by at least one additional item. This result was pursued recently by Giesbrecht and Di Lollo (1996, 1998). These authors compared a blank interval following the second target with three other conditions: the second target was followed by the standard stream of items in the RSVP procedure, the second target was followed by a single item (interruption mask), and the second target was superimposed on another pattern (integration masking). Unlike the finding for first-target masking (Seiffert & Di Lollo, 1997), the AB was observed only in the standard stream and interruption masking conditions. Furthermore, a single delayed masking item was as effective as an entire stream of items in producing the AB. Therefore, only interruption masking of the second target was effective in producing the AB.

Giesbrecht and Di Lollo (1998) argued that this asymmetry of masking points to two distinct roles for mask-

ing. Masking of the first target is merely one of potentially many ways to place stress on the attentional processes involved in target identification so that an AB can be observed. However, masking of the second target plays a much more specific role. In this case, masking by interruption permits the representation of the second target to be substituted in consciousness for the masking item while attention is occupied with the first target item. Di Lollo and colleagues (Di Lollo, Enns, & Rensink, in press; Enns & Di Lollo, 1997; Giesbrecht & Di Lollo, 1998) prefer the term *substitution*, rather than *interruption*, for the masking involved in second-target processing, because the mask item does more than simply terminate processing of the target. The mask itself appears to be the new focus of object recognition mechanisms.

Object substitution occurs because perception is based on an ongoing comparison between provisional high-level percepts and the low-level activity triggered by current sensory input (Di Lollo et al., in press). When attention is switched to a new location or object, the perceptual cycle begins with the sensory input at that location or object. In AB experiments, when attention switches from the first to the second target item, the available sensory input will more likely come from the item that trails the second target than from the target item itself, since attention switching itself requires some time. This substitution process can be observed quite directly in the false-positive reports that observers make for the items that follow and, therefore, mask the second target (Chun, 1997; Martin, Isaak, & Shapiro, 1995).

Scope of the Present Study

In the experiments reported here, two types of masking were combined with each of two targets in an orthogonal design. Integration masking (simultaneously presented visual noise) and interruption masking (a briefly delayed visual pattern in the target location) were chosen because of their differential sensitivity to manipulations of attention.

Since previous experiments have pointed to the possibility of two separate roles for masking in the AB, it is tempting to simply assume that the effects are independent. Namely, masking of the first target is one way to place a heavy demand on the attentional system, whereas masking of the second target is critical for obtaining an object substitution effect. However, this conclusion is unwarranted without a formal test—hence, the present set of experiments. Finding no interaction between the effects of masking the first target and those of masking the second target would lend strong support for the hypothesis that masking plays two roles. Specifically, the form of masking should not matter for the first target, whereas an interruption mask following the second target should be critical to the AB. Alternatively, finding clear interactions between masking effects on the two targets would force a reconsideration of the hypothesis.

A secondary motivation was to extend what had already been learned about masking and AB, using the RSVP procedure, to the simplified Duncan et al. (1994) methodology. A recent paper (Ward et al., 1997) has shown that these methods yield similar results in many respects, although the role of masking was not investigated systematically for either target in that paper.

To anticipate our results, in Experiment 1, we found strong evidence favoring independent roles of masking for the first target and the second target. Applying either mask to the first target was effective in reducing accuracy in the second target. However, only the interruption mask was effective in doing so when applied to the second target. Further evidence for the hypothesis of independence was obtained in Experiment 2, in which the overall level of integration masking was varied systematically. The pattern of accuracy obtained for integration masks was shown to be similar at all levels of accuracy. Finally, the case for independence was strengthened in Experiment 3, in which we mixed aspects of integration and interruption masking for each of the targets and observed predictable consequences of such mixing.

EXPERIMENT 1

The methods used in each experiment followed those of Duncan et al. (1994), in which observers were asked to report the identity of two letter targets, briefly presented and then masked by digits, at two of the four cardinal locations centered around the fixation point. The temporal interval between the presentation of the first and second targets was randomly varied in 100-msec steps within a range of 0–600 msec. Identification accuracy was examined as a function of two main factors: the targets (first, second) and the kind of mask (integration vs. interruption). The integration mask consisted of a digit spatiotemporally superimposed on the target letter. The interruption mask consisted of the same digit, delayed by 90 msec, in the target location.

Method

Observers. Twenty-four undergraduate volunteers (19 female, 5 male; age range = 17–25 years) participated for either extra course credit or \$5. The observers in this and all subsequent experiments reported normal or corrected-to-normal vision, and all were naive to the purpose of the experiment.

Stimuli and Apparatus. The target items were chosen randomly from 20 uppercase alphabetic characters similar to Letraset No. 8831. The letters I, O, P, Q, S, and Z were excluded because of confusability with the masking digits. The masking items were chosen from the digits 0–9, with the 6 being excluded because it was confusable with several letters.

The items were plotted on a Tektronix 608 oscilloscopic point-plotter equipped with fast P15 phosphor. At a viewing distance of 57 cm, set by a headrest, all characters were approximately 1.0° in height. Screen illumination was 55 cd/m², provided by a pair of neutral-density filtered fluorescent tubes located to the front and sides of the oscilloscope screen. All targets and masks were pre-

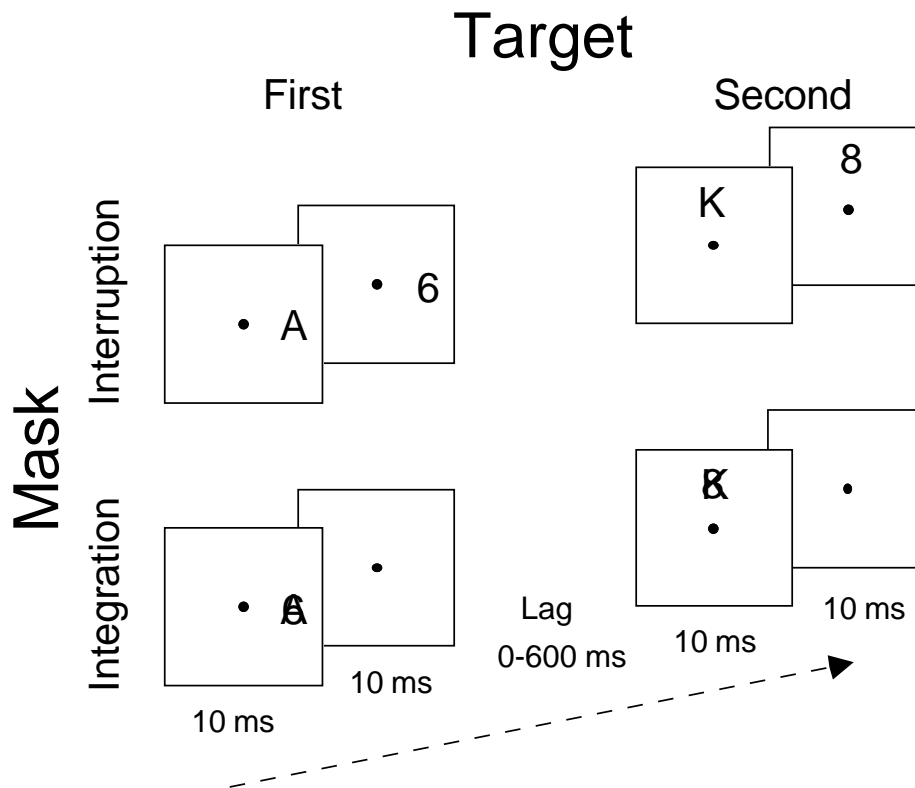


Figure 1. Illustration of the stimulus displays in Experiment 1. Letter targets were displayed in two of four possible locations. Each target (first, second) was accompanied either by an interruption mask, which followed the target after 90 msec, or by an integration mask, which was spatiotemporally superimposed on the target. The temporal lag between the first and second targets varied between 0 and 600 msec.

sented for 10 msec at a luminance of 500 cd/m², as measured by a Minolta LS-100 photometer. A fixation point was continuously present at the center of the 8° × 8° screen. The first target and its mask were presented randomly 1° above, below, left, or right of fixation. The second target and mask appeared randomly in one of the three remaining locations, as is shown in Figure 1.

Procedure. Testing was done individually in a dark room. A small fluorescent tube illuminated the keyboard. The experimenter was in an adjacent room, communicating with the observer over a loudspeaker.

The observers were required to fixate a central fixation point throughout each trial. They initiated each trial by pressing the space bar on the keyboard. After a random interval of 0–500 msec, the first target (T1) was presented for 10 msec in one of the four possible locations. This was followed, after a variable interval (0–600 msec) by the second target (T2), also for 10 msec, in one of the remaining three locations. The observers were asked to identify both letters by depressing the corresponding keys. They were warned that the letters I, O, P, Q, S, and Z would never appear, that the two target letters would never be identical on a given trial, and that the order of letter identification did not matter.

Figure 1 illustrates the possible combinations of trial sequences: Each of the two targets could be masked by either integration or interruption. These conditions were denoted *integration–integration*, *integration–interruption*, *interruption–interruption*, and *interruption–integration*, where the masking labels referred to the type of

mask used for the first and second targets, respectively. In the integration masking conditions, the masking digit was spatiotemporally superimposed on the target letter. In the interruption masking conditions, the digits were presented in the same location but after a delay of 80 msec. This meant that the total SOA between targets and masks was 90 msec.

The observers were given 10 practice trials prior to testing. Testing consisted of four blocks of 140 trials, with each block lasting approximately 8–10 min. Only one of the four combinations of target and mask was tested within a single block, with the order of the four conditions counterbalanced over blocks.

Results

Accuracy of target identification is shown in Figure 2, separately for the two targets in each of the four masking conditions. The main result was that the characteristic features of the AB—a second-target deficit accompanied by a U-shaped pattern of accuracy over lag for the second target—were evident only when both targets were masked by interruption (Figure 2A). In stark contrast to this pattern of accuracy, when both targets were masked by integration (Figure 2B), there was neither a selective second-target deficit in accuracy nor a U-shaped pattern of accuracy over lag. This clearly indicated that the type of

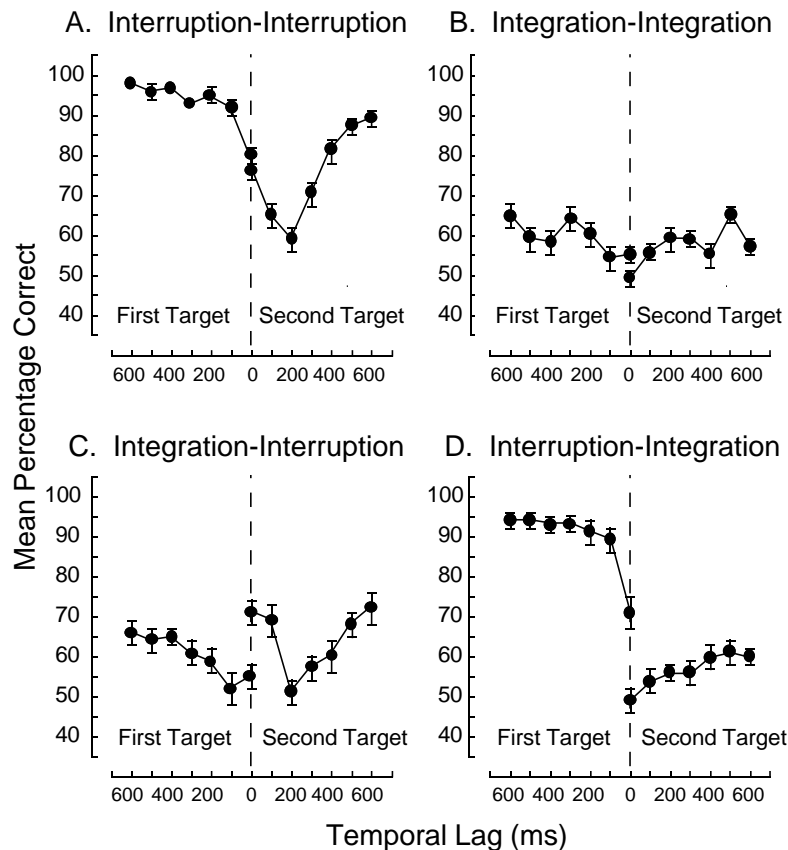


Figure 2. Mean accuracy of target identification in Experiment 1. (A) Interruption-Interruption: First and second targets were followed by interruption masks. (B) Integration-Integration: First and second targets were embedded in integration masks. (C) Integration-Interruption: First target was embedded in an integration mask, and second target was followed by an interruption mask. (D) Interruption-Integration: First target was followed by an interruption mask, and second target was embedded in an integration mask. Error bars represent standard errors of the mean. Chance accuracy level was 5%.

masking had a strong influence on target accuracy in this experiment.

The kind of influence exerted by masking could be seen in greater detail by examining the conditions in which a different mask was used for the first and second targets. When the first target was masked by integration and the second by interruption (Figure 2C), the resulting pattern of accuracy for the first target was similar to the first-target integration condition in Figure 2B; the second-target pattern was similar to the second-target interruption condition in Figure 2A. The same principle of independence applied when the masking types were reversed (Figure 2D). Here, the accuracy pattern for the first target masked by interruption was similar to the first-target interruption condition in Figure 2A and the pattern for the second target masked by integration was similar to the second-target integration condition in Figure 2B.

Several other effects, all secondary to these main findings, were that first-target accuracy was impaired much

more by integration than by interruption masking (left panels of Figures 2B and 2C vs. Figures 2A and 2D), second-target accuracy was approximately equal when the two forms of masking were at their maximum (left vs. right panels of Figure 2C), and accuracy with integration masking increased monotonically for both targets with an increase in temporal lag (Figures 2B, 2C, and 2D). These observations were supported by the following statistical analyses.

The statistical analyses consisted of a repeated measures analysis of variance (ANOVA) in which accuracy was examined as a function of target (first, second), first-target mask (integration, interruption), second-target mask (integration, interruption), and lag (0, 100, 200, 300, 400, 500, 600, 700 msec). The data were also analyzed using a contingent measure of target accuracy (i.e., accuracy for a given target contingent on the other target being accurately identified on the same trial). Since the results of these data did not differ in any statistical detail

from the analysis of raw accuracy, we chose to report the raw accuracy.

All the effects were statistically significant in the ANOVA, and so we followed up with polynomial contrast analyses of the effects of lag. The U-shaped pattern of accuracy over lag reveals itself in these analyses as a significant quadratic component. The only lag effects that contained a significant quadratic component were observed when the second target was masked by interruption [right panel of Figure 2A, $F(1,6) = 31.49$, $MS_e = 148.94$, $p < .001$; right panel of Figure 2C, $F(1,6) = 37.83$, $MS_e = 177.58$, $p < .001$; combined, $F(1,6) = 61.056$, $MS_e = 185.36$, $p < .001$].

The other significant lag effects consisted of linear components, indicating monotonic improvements with lag for first-target accuracy [integration masking, $F(1,6) = 29.33$, $MS_e = 156.66$, $p < .001$; interruption masking, $F(1,6) = 163.49$, $MS_e = 50.62$, $p < .001$] and monotonic improvements with lag for second-target accuracy [integration masking, $F(1,6) = 27.57$, $MS_e = 119.60$, $p < .001$].

Inspection of the three-way interaction of target \times first-target mask \times second-target mask [$F(1,23) = 5.12$, $MS_e = 48.28$, $p < .03$] indicated that accuracy was significantly higher on the first target than the second target only when first target was followed by an interruption mask (Figures 2A and 2D) [$F(1,23) = 175.83$, $p < .001$]. Target accuracy did not differ significantly when the first target was masked by integration (Figures 2B and 2C) ($F < 1$).

A simple-effects ANOVA for first-target accuracy revealed a significant effect of mask [$F(1,23) = 318.77$, $MS_e = 73.18$, $p < .001$]: Accuracy was much lower when the first target was masked by integration (59.81%) than when it was masked by interruption (90.91%). There was also a main effect of the second-target mask [$F(1,23) = 8.81$, $MS_e = 16.43$, $p = .007$]: First-target accuracy was lower when the second target was masked by integration (74.17%) than when it was masked by interruption (76.62%). The interaction between mask types and the two targets was not significant ($F < 1$).

The same analysis for the second target also revealed a significant effect of mask [$F(1,23) = 31.85$, $MS_e = 125.71$, $p < .001$]: Second-target accuracy was lower when the second target was masked by integration (56.79%) than when it was masked by interruption (69.70%). The analysis also revealed a main effect of the first-target mask [$F(1,23) = 11.21$, $MS_e = 60.09$, $p = .003$]: Second-target accuracy was lower when the first target was masked by integration (60.56%) than when it was masked by interruption (65.89%). Finally, only for second-target accuracy was there a significant interaction between mask types and their effects on the two targets [$F(1,23) = 19.19$, $MS_e = 45.2$, $p = .001$], reflecting that increases in the difficulty of first-target identification (through integration masking) was especially detrimental to accuracy on the second target.

Discussion

The results of Experiment 1 indicate that masking plays two independent roles in the AB. Consider first the accuracy of report for the first target. Here, integration masking was more effective in impairing target accuracy than was interruption masking. Yet, this difference in masking of the first target was unaffected by the kind of masking that the second target was subjected to on the same trial (compare Figures 2A and 2C with Figures 2B and 2D). Furthermore, both types of mask applied to the first target were sufficient to produce U-shaped patterns of accuracy over lag for the second target (Figures 2A and 2C).

With regard to accuracy in reporting the second target, interruption masks produced a U-shaped pattern of accuracy over lag, whereas integration masking produced a shallow monotonic trend of improvement in accuracy with lag. As with the first-target masking effects, this difference in second-target masking effects was uninfluenced by the kind of mask used in conjunction with the first target (compare Figures 2A and 2B with Figures 2C and 2D).

These results are therefore not consistent with the view that the primary role played by visual masking in the AB is to bring second-target accuracy into a measurable range. This was accomplished in this experiment by using both integration and interruption masks with the first target. Nevertheless, the U-shaped pattern of accuracy for second-target identification was observed only when the second target was masked by interruption. When both targets were masked by integration, not only was there no U-shaped pattern of second target accuracy but there was no second-target deficit of any kind (Figure 2B). Accuracy on the second target was equivalent to that on the first target.

Distinguishing an Attentional Blink From a Two-Object Cost

What these results make explicit is that several features of the typical AB pattern of data are really quite separable from one another. These include a second-target deficit in accuracy (i.e., lower accuracy on the second target than on the first target), a U-shaped pattern of accuracy over temporal lag, and the average or baseline level of response accuracy. The results of Experiment 1 show that a second-target deficit is not always observed (Figures 2B and 2C), that a U-shaped pattern of accuracy over lag is not a universal feature of second-target accuracy (Figures 2B and 2D), and that similar U-shaped patterns can be observed for second-target accuracy in conjunction with very different levels of accuracy for the first target (Figures 2A and 2C).

The most striking result in this regard occurred when the first and second targets were each masked by integration (Figure 2B). Here, despite baseline accuracy for both targets being in a measurable range (despite chance accuracy being 5%, average accuracy was over 50%),

there was no evidence of either a selective second-target deficit or a U-shaped pattern of accuracy over lag for the second target. Indeed, the lag effects for both targets were statistically symmetrical, with small improvements in target accuracy with increasing lag. Thus, there was no “attentional blink” by conventional operations.

Yet, the symmetrical effects of lag in this condition point to an impairment in performance that is a function of the temporal proximity of the two targets. The data in Figure 2 show that this monotonic lag effect, which is present in all conditions, is experimentally separable from the U-shaped pattern of accuracy that occurs only for second targets and then only when they are masked by interruption. We interpret the monotonic lag effect as evidence of a “two-object cost” (Duncan, 1984). That is, the processes involved in identifying a target are limited in capacity; if two targets are presented for identification either simultaneously or in rapid succession, then the accuracy and speed of identification of both will be impaired. Note that there is nothing in this account that predicts that accuracy for either the first target or the second target will be impaired selectively. We therefore conclude that the AB (i.e., the second-target deficit) does not arise from a two-object cost per se but rather arises from the object substitution effect produced by the mask that follows the second target in time.

It is also worth noting that, in some combinations of mask types for the two targets, a second-target deficit occurred simply because of the relative levels of mask strength used in each case (Figure 2D). As will be seen in Experiment 2, this is simply an accident of combining masks at different levels of strength. In this case, the interruption mask is not nearly as effective as the integration mask in impairing target accuracy. This should not be mistaken for a second-target deficit that occurs when masking strength is equated.

EXPERIMENT 2

One concern in comparing the effects of integration and interruption masks in Experiment 1 was that integration masking yielded generally lower accuracy in all conditions. It is possible that the relatively flat accuracy functions observed for the integration masks were an unintended by-product of reduced accuracy levels.

In Experiment 2, we examined integration masking more closely for both targets by testing three different mask luminance levels. These levels were intended to span a wide range of task difficulty, from high accuracy (a dim mask) to accuracy as low as in Experiment 1 (a bright mask). If the absence of U-shaped accuracy functions for the second target was the result of the strong form of integration masking we had used, then decreasing task difficulty by reducing mask luminance should reinstate the effect. However, if integration masking of the second target presents the visual system with a fundamentally different problem than interruption masking, then




Figure 3. Illustration of the stimulus displays in Experiment 2. Both letter targets were embedded in integration masks. The masking digit was dim (200 cd/m²), medium (300 cd/m²), or bright (500 cd/m²).

the same pattern of accuracy should merely be seen at different baseline levels of performance.

Method

Twenty-four observers (19 female, 5 male; age range = 17–25 years) participated for extra course credit in an integration–integration condition (Figures 1 and 2B). All methodological details were the same as in Experiment 1, with the exception that three different mask luminance levels were tested, as shown in Figure 3. The bright mask (500 cd/m²) was identical to that in Experiment 1, in which target and mask luminances were equal. The medium mask (300 cd/m²) and dim mask (200 cd/m²) conditions were designed to improve target identification accuracy. Testing was done in three counter-balanced blocks of 140 trials, with one block devoted to each mask luminance level.

Results

Accuracy of target identification is shown in Figure 4. The main result was that mask luminance had no influence on the shape of the accuracy functions over temporal lag. Although variations in mask luminance were effective in changing overall levels of accuracy, as we had intended, accuracy improved slightly with increases in lag and in an equivalent way for both targets. These observations were supported by statistical analyses.

A repeated measures ANOVA revealed that accuracy decreased significantly with mask luminance [$F(2,46) = 124.49$, $MS_e = 445.74$, $p < .001$] and that accuracy in-

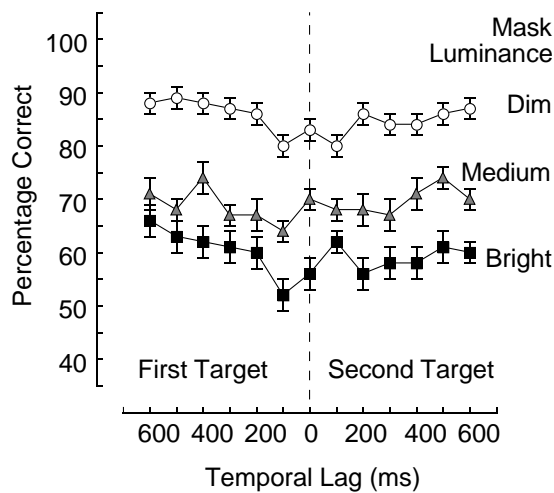


Figure 4. Mean accuracy of target identification in Experiment 2. Error bars represent standard errors of the mean. Chance accuracy level was 5%.

creased significantly with temporal lag [$F(6,138) = 6.12$, $MS_e = 117.11$, $p < .001$]. Accuracy did not vary significantly with target [$F(1,23) = 2.92$, $MS_e = 154.54$, $p > .10$]. The only significant interaction involved target \times lag [$F(6,138) = 4.08$, $MS_e = 72.92$, $p < .001$], which reflected larger accuracy improvements with lag for the first target than for the second target. Polynomial contrast analyses of the effects of lag revealed significant linear components for both the first target [$F(6,138) = 23.35$, $MS_e = 107.05$, $p < .001$] and the second target [$F(6,138) = 13.81$, $MS_e = 82.98$, $p = .001$]. No higher order components were significant ($F < 1$, for all quadratic components).

Discussion

These results confirm that integration masking of the second target in the AB methodology presents the visual system with a fundamentally different problem than interruption masking. Despite testing over a wide range of task difficulty, integration masking of the second target never yielded a U-shaped pattern of accuracy over temporal lag, and it never yielded a second-target deficit. These are both features of the data that are readily obtained with interruption masking of the second target.

EXPERIMENT 3

As noted in the introduction, the time courses of integration and interruption masking are considerably different. The effects of an integration mask are optimal when the mask is presented simultaneously with the target (0-msec SOA), whereas interruption masks are at their strongest at some positive-valued SOA. In Experiments 1 and 2, we employed a 0-msec SOA condition in order to maximize integration masking and a 90-msec SOA to maximize interruption masking.

In Experiment 3, we explored the consequences of combining the effects of integration and interruption masking on target identification. We accomplished this by testing an SOA condition (60 msec) that was intermediate to those tested previously. If the type of masking with respect to the first target is really immaterial to the AB, as we have argued, then the combined effects of masking in this condition should have little effect on first-target identification other than to reduce accuracy by an amount equivalent to the visual noise accompanying the integration mask.

If the type of mask with respect to the second target is critical, as we have also argued, then a different set of predictions holds for the second target in the intermediate SOA (60-msec) condition. Specifically, the combined effects of integration and interruption masking should have three testable consequences. First, overall levels of accuracy for T2 should be reduced by the extent of the integration masking at an SOA of 60 msec. Second, there should still be evidence of a U-shaped pattern of accuracy for T2, reflecting the component of interruption masking in this condition. Finally, the degree of recovery of T2 accuracy as a function of lag should be noticeably reduced relative to the 90-msec condition, reflecting the increased amount of integration masking.

Method

Twenty-four different observers (17 female, 7 male; age range = 18–24 years) participated for extra course credit. The observers were tested in three different target–mask SOA conditions (0, 60, or 90 msec). In each, the SOA corresponded to the interval between the onset of each target and the mask. Testing was done in three counterbalanced blocks of 140 trials, with one block devoted to each SOA condition. All other methodological details were identical to those in Experiment 1.




Figure 5. Mean accuracy of target identification in Experiment 3. Error bars represent standard errors of the mean. Chance accuracy level was 5%.

Results

Accuracy of target identification is shown in Figure 5. The main result was that variations in target-mask SOA affected primarily the overall level of accuracy for the first target. However, these same variations in SOA resulted in qualitatively different accuracy functions for the second target. These observations were supported by statistical analyses.

A repeated measures ANOVA showed that all effects were significant, and so the subsequent analyses focused on patterns of accuracy over lag for each of the targets. A simple effects analysis of T1 indicated that accuracy increased significantly with increases in the lag separating the two targets [$F(6,138) = 42.27$, $MS_e = 71.79$, $p < .001$], that accuracy increased significantly with an increase in SOA [$F(2,46) = 166.52$, $MS_e = 339.09$, $p < .001$], and that the interaction between lag and SOA was significant [$F(12,276) = 8.902$, $MS_e = 64.80$, $p < .001$]. This interaction was found to reflect smaller SOA differences only at lags of 0 and 100 msec, reflecting greater interference from temporally near targets when masking was by interruption (SOA = 60 and 90 msec). When lags 0 and 100 msec were removed from the analysis, the interaction was no longer significant [$F(8,184) = 1.29$, $MS_e = 57.73$, $p > .20$].

The same analysis for T2 revealed significant main effects of lag [$F(6,138) = 15.54$, $MS_e = 117.97$, $p < .001$] and SOA [$F(2,46) = 39.35$, $MS_e = 556.85$, $p < .001$], along with a significant interaction [$F(12,276) = 10.32$, $MS_e = 136.11$, $p < .001$]. Polynomial contrast analyses revealed significant quadratic components in all three SOA conditions [0 msec, $F(6,23) = 6.27$, $MS_e = 111.0$, $p = .013$; 60 msec, $F(6,23) = 20.72$, $MS_e = 169.39$, $p < .001$; 90 msec, $F(6,23) = 88.26$, $MS_e = 109.80$, $p < .001$]. However, the sign of the quadratic component in the 0-msec condition, in which accuracy increased slightly in the intermediate lags, was the opposite of that in the 60- and 90-msec conditions, in which it decreased significantly before beginning to improve at longer SOAs. In summary, the characteristic U-shaped pattern of accuracy for the second target was observed only in the 60- and 90-msec SOA conditions.

The degree of recovery in second-target accuracy with increasing lag was compared in the 60- and 90-msec SOA conditions by examining the estimates of the linear slope over temporal lags from 200 msec (maximum masking) to 600 msec (least masking). The mean slope for the 90-msec condition (8.3% for each additional 100 msec of lag) was significantly larger than that for the 60-msec condition (4.3% for each additional 100 msec of lag) [$F(1,46) = 16.42$, $MS_e = 1,194.17$, $p < .001$], indicating less recovery in the 60-msec condition.

Discussion

These results clearly indicate that masking plays a different role for the first target in the AB task than it does for the second target. Combining aspects of integration and interruption masking had little effect on identifica-

tion of the first target other than affecting a general decrease in accuracy. In contrast, mixing the two forms of masking had predictable consequences on the shape of the accuracy function for the second target. Not only was accuracy generally depressed in the 60-msec condition, relative to the 90-msec condition, but the two other predictions were confirmed. There was a measurable U-shaped pattern of accuracy in the 60-msec condition, consistent with some degree of interruption masking. The degree of recovery in accuracy over lag was also considerably less in the 60-msec condition than in the 90-msec condition, reflecting the integration masking component. Combining integration and interruption masking is therefore another way to dissociate the two roles played by visual masking in the AB.

GENERAL DISCUSSION

The experiments reported here each point to two distinct roles for visual masking in the AB. A systematic comparison of two kinds of masking (integration and interruption) in Experiment 1 revealed that, for the first target, either type of masking was sufficient to produce the signature of the AB: a U-shaped pattern of second-target accuracy over temporal lag. However, masking requirements were much more specific for the second target. Here, only interruption masking of the second target produced the U-shaped pattern of accuracy.

This conclusion was shown to be robust in two ways in subsequent experiments. Experiment 2 demonstrated that the ineffectiveness of integration masking with respect to the second-target deficit was not contingent on lower overall levels of accuracy. Integration masking of the second target did not result in a U-shaped pattern of accuracy over temporal lag, or even a more general second-target deficit, at any mask intensity. Experiment 3 revealed a dissociation in masking effects for first and second targets in another way. Whereas combining the two forms of masking had no differential effect on the accuracy pattern over temporal lag for the first target, it had predictable consequences for the second target. There, it produced qualitative differences in the shape of the accuracy function over temporal lag.

Theoretical Considerations

There have been three main theories proposed to date to account for the AB. In what follows, we briefly summarize each theory, indicating those aspects of the present data that are consistent and inconsistent with each view. Finally, we summarize our emerging view of the critical importance of object substitution masking to the AB.

Visual short-term memory (VSTM) theory. The first theory of the AB is premised on a competition among rapidly presented items within a limited-capacity VSTM (Shapiro et al., 1994). Early visual representations of items in a temporal stream are believed to gain access to the more sophisticated operations of VSTM, largely on the basis of the degree to which they match preset templates.

However, control over access is not perfect, such that non-matching items may gain entry if they are temporally contiguous to the target items. Once in VSTM, items compete for limited processing resources that are allocated to them jointly on the basis of the order of entry into VSTM (first-come–first-served) and on the basis of their similarity to the preset templates.

The VSTM theory accounts for the main features of the AB as follows. Both the first-target and the second-target items match the preset templates; therefore, their probability of admission to VSTM is high. Furthermore, items immediately following the targets (i.e., masks) are also likely to be admitted because of their temporal contiguity to targets. Once in VSTM, the first target will receive the majority of processing resources, because it both matches a template and is the first item to be entered. The item immediately following the first target is also likely to be admitted, but it will be allocated fewer resources. If that item happens to be the second target, it will fare quite well in the competition with other nontarget items. However, if it does not match any templates (i.e., it is a masking item), then it will suffer. In general, this account predicts second-target deficits in accuracy at intermediate temporal lags, since it is at these lags that limited resources will already have been allocated to the first target, and other nontarget stream items will have consumed the remaining resources. At longer lags, resources will have become available again, once the processing of the first target is complete.

How does the VSTM theory fare in accounting for the finding that visual masking plays two different roles in the AB? Although the theory was not intended to address itself to masking, the role played by masking in the identification of the first target follows quite naturally from it. Consider first that, with no masking of the first target, there is much less of a competition in VSTM for limited resources, and so target identification can proceed quite quickly. With a mask, either one of integration or one of interruption, the load on VSTM will be increased substantially, and so the time required to identify the first target will be increased. Indeed, this is the account that has been given to explain the reduction in the AB that occurs when observers are tested with a stream that is missing an item immediately following the second target (Raymond et al., 1992; Seiffert & Di Lollo, 1997).

In contrast, the VSTM theory copes much less well with the finding that interruption masking of the second target is essential to the AB. In fact, this theory does not make any predictions that hinge on the nature of the items that accompany or follow the second target. There is only the same general prediction as there is for the first target: Forcing the visual system to individuate the second target from other nontarget items should serve to reduce overall accuracy for that target. There is no provision for mask-dependent differences in the temporal lag function.

If anything, VSTM theory might be expected to predict masking effects that are opposite in direction to those

found in the present study. Consider that admission of a nontarget item into VSTM depends on temporal contiguity to items that match the templates (i.e., the targets). This predicts that nontarget items presented simultaneously with the target (integration masking) would be more likely to be admitted into VSTM than those that follow the target after a 90-msec delay (interruption masking). However, in direct contradiction to this prediction, the present findings clearly showed that interruption masking of the second target produced the larger deficit in identification. VSTM theory, therefore, fails to account for the necessity of following the second target with an interruption mask.

Dwell-time theory. A second theory proposed to explain the AB focuses on the amount of time required by limited-capacity processes to identify an item (Duncan et al., 1994). The underlying premise is that the natural unit of visual attention is the “object.” The multiple attributes of the same object (e.g., its shape, color, motion) are processed in parallel, whereas the attributes from separate objects are processed serially. *Dwell time* refers to the amount of time required for the parallel pathways to complete their operation for a single object. In experiments involving spatial arrays of multiple objects, a two-object cost can be measured when observers are asked to report the attributes, or identify more than one, of the objects (Duncan, 1980, 1984, 1993). In experiments involving temporal streams of visual items, a similar two-object cost is reflected in the AB or dwell-time measure.

According to dwell-time theory, when the first target is detected in the visual stream, attention is devoted to its processing. When the second target then appears before processing of the first is complete, there is a reduction in accuracy for the second target, reflecting the dwell time associated with the first target. Visual masking plays a role in this theory only insofar as it brings performance accuracy into a range where the dwell-time effect is measurable. The specific type of masking should be irrelevant.

How does dwell-time theory fare with the present results? The one finding of this study that is entirely consistent with dwell-time theory is the general improvement in accuracy that was seen for both targets with increasing lag. The significant linear trends over lag associated with integration masking in all three experiments are consistent with there being a cost associated with processing two targets in close temporal proximity. However, it is also important to note that this two-object cost reflected in the temporal lag effects was quite small in comparison with the lag effects associated with interruption masking of the second target. Dwell-time theory makes no predictions concerning the critical importance of interruption masking of the second target.

Two-stage theory. A third account of the AB begins with the premise that visual processing is composed of two distinct and sequential stages (Chun & Potter, 1995). In the first stage, the image is registered rapidly in terms of the fundamental features of the visual system (e.g., oriented blobs, colors, motion direction), and each item

is identified provisionally, presumably by matching these features with templates that are appropriate for the task. This representation is short-lived if the visual display is brief, but some categorical selection is possible, on the basis of the briefly available identity. Items thus selected are admitted into a second, attention-limited, stage of processing, where more sophisticated analyses of recognition and response planning may be carried out. However, the increased sophistication of the second stage comes with a cost, in that operations are much slower and can be performed on only one item at a time.

The AB occurs, in this theory, because the attentive stage of processing has not been completed for the first target in the visual stream before the second target is briefly presented. The short-lived preattentive representation of the second target is rendered useless for further processing by two different operations. First, the representation decays rapidly over time. Therefore, if too much time has elapsed while the first target is being processed by the attentive stage, there will be no trace of the second target for similar processing. Second, if nontarget items follow the second target before it can be admitted, then its preattentive representation will be overwritten before it can receive attentive processing. In either case, an accuracy impairment will be the result for the second target.

Like the other two theories, the importance of masking of the first target is predicted by this theory. If anything, two-stage theory is the most explicit in establishing a link between the difficulty of processing the first target and the extent and magnitude of the AB for the second target. According to two-stage theory, any factor that contributes to slower identification of the first target should lead to a larger AB. In this view, integration and interruption masking simply happen to be two convenient ways by which this can be accomplished.

Two-stage theory is unique, however, in providing a differential prediction with regard to integration versus interruption masking of the second target. Whereas integration masking of the second target will lead to impairments in accuracy through the passive process of decay, interruption masking will bring with it the potential for overwriting of the second target at the preattentive level by a subsequent masking item. Such overwriting might even be expected to produce a more catastrophic impairment in accuracy, since the trace of the second target may vanish abruptly at the time of overwriting. By contrast, the limiting factor in integration masking would be the duration of the fading trace of the second target, a process that is by definition more gradual.

Object Substitution

Although we find ourselves in general agreement with the two-stage theory (Chun & Potter, 1995), we prefer the concept of *substitution* to that of interruption or even overwriting, to describe the role played by the delayed mask with respect to the second target in the AB (Enns & Di Lollo, 1997; Giesbrecht & Di Lollo, 1998). Whereas interruption places emphasis on the termination of pro-

cessing for a visual item, and overwriting leaves open the possibility that the representation of a newer item is simply superimposed on the decaying trace of the older item, substitution is intended to convey the sense that the original item has been completely replaced by the temporally trailing item, at least as far as conscious processes are concerned. Direct evidence that the item following the second target does more than simply terminate processing can be seen in the large number of false-positive reports that observers make for the item immediately following the second target in an RSVP stream (Chun, 1997; Martin et al., 1995). Evidence that the representation of a delayed mask is not merely superimposed on that of the target is seen in the large interruption masking effects that can be obtained using only four small dots that surround, but that do not touch, a target item (Enns & Di Lollo, 1997).

At the heart of the object substitution hypothesis is the view that visual representations of attended items are fundamentally different from those of unattended items. Of particular importance to the perception of rapid temporal sequences and visual masking is the increased spatiotemporal resolution and durability of attended items (Enns, Brehaut, & Shore, 1996; Moran & Desimone, 1985; Posner, 1980; Rensink, O'Regan, & Clark, 1997; Treisman & Gelade, 1980; Tsal, Meiran, & Lamy, 1995). The lower resolution and increased volatility of unattended items therefore leaves them more vulnerable to substitution by masking items.

Two specific predictions ensue for the relations between interruption masking and focused attention. The first is that an interruption mask will do very little to reduce the visibility of an item that is the current focus of attention. The second is that an interruption mask is likely to replace the trace of the target item when that item and the mask are presented outside the focus of attention. These predictions can be applied equally well to experiments involving multiple items distributed over space or over time. For example, if multiple items are arrayed over space, and one of the items is unpredictably masked shortly thereafter, it is the mask (and not the target item) that is available for conscious report. In contrast, if the item to be masked is cued in advance, or if the item is presented alone in the visual field, then the same mask will have little effect on target accuracy. Both of these predictions have been confirmed in recent experiments (Di Lollo et al., in press; Enns & Di Lollo, 1997).

The corresponding predictions for multiple items distributed over time can be tested directly in the present study. One assumption that must be made explicit in order to sustain this analogy is that the first target in the temporal stream is attended more quickly and efficiently than the second target. There are any number of reasons to justify this assumption, including that the first target is presented when the observer is in an optimal state of readiness and that identification of the first target uses limited resources that are required in the identification of the second target. In any case, the important prediction

is that an interruption mask should have little effect on the first target, which is readily attended, whereas the same mask should have devastating effects on accuracy for the second target, which can be attended less readily. This is exactly the pattern of results shown in Figures 2A and 5 for an interruption mask.

However, there also appears to be something of a paradox in both the predictions that follow from the object substitution account and the results of the present experiments. In both cases, an interruption mask following the first target is said to be ineffective in reducing the accuracy of first-target identification; yet, this same mask is said to be critical in reducing second-target accuracy by postponing its identification. Could an interruption mask actually have its primary influence not on the item that immediately preceded it (the first target) but on an item that is not at the same spatial location and occurs later in time (the second target)? Evidence in direct support of this hypothesis has been reported recently in three independent studies (Arnell & Duncan, 1997; Jolicoeur & Dell'Acqua, 1998; Zuvic, Visser, & Di Lollo, in press). In each study, observers made speeded responses to a non-masked second target presented at various lags following a first target, which was masked by a trailing pattern (interruption). As in the present study, identification of the first target was unaffected by the delayed mask. Nevertheless, masking of the first target had a pronounced effect on response times to the second target: Responses were longest when the mask was presented directly after the first target and became shorter as the lag was increased. In contrast, when the first target was not masked, response times to the second target were shorter and remained constant across all lags. This is a clear indication that interruption masking of the first target slows the processing of the second target. It is during this delay that the low-level representation of the second target remains vulnerable to masking and the AB is produced.

Equally important in confirming the object substitution hypothesis in the present data is the prediction that another form of masking (one not involving substitution) should produce a different pattern of results. This was confirmed in the conditions involving exclusively integration masking (Figures 2B, 4, and 5). Here, there was no detectable asymmetry in masking effects for the first and second targets. The only effect was that of temporal lag. Placing two targets in temporal proximity to one another resulted in a small impairment in accuracy for both targets, consistent with a two-object cost in identification (Duncan, 1980, 1984, 1993).

Conclusion

These experiments have demonstrated that masking plays two separate roles in the AB. Masking of the first target, whether by integration or interruption, introduces a delay in processing the second target. During this delay, the second target is vulnerable to being replaced by a trailing mask with a probability that diminishes as the temporal lag between targets is increased. Note that, for

this role, visual masking per se is not critical. Any number of other ways to prolong the identification process should serve as well (e.g., Chun, 1997; Chun & Potter, 1995). However, it is of both methodological and theoretical importance to acknowledge that this is the role played by the item immediately following the first target in the conventional RSVP stream.

The second, and most important, role played by visual masking concerns the second target. Here, we have shown that interruption masking is essential in obtaining the second-target deficit in accuracy. Other ways of decreasing accuracy for the second target (e.g., integration masking) simply do not produce the same result. Although this finding is also of methodological interest, we believe it is of even greater theoretical importance, since it enables a better understanding of the hidden visual processes that normally occur outside the focus of attention and consciousness. One of these processes appears to be the ready substitution of visual representations by those of temporally trailing items when the original representations are not attended. We believe it is this process that underlies the recently reported demonstrations that major changes to a visual display can go unnoticed, unless attention is deliberately deployed to the region of the changing element (Rensink et al., 1997; Simons, 1996).

REFERENCES

- ARNELL, K., & DUNCAN, J. (1997, November). *Speeded responses or masking can produce an attentional blink*. Poster presented at the annual meeting of the Psychonomic Society, Philadelphia.
- BACHMANN, T., & ALLIK, J. (1976). Integration and interruption in the masking of form by form. *Perception*, *5*, 79-97.
- BREITMEYER, B. G. (1984). *Visual masking: An integrative approach*. New York: Oxford University Press.
- CHUN, M. M. (1997). Temporal binding errors are redistributed by the attentional blink. *Perception & Psychophysics*, *59*, 1191-1199.
- 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 & Performance*, *21*, 109-127.
- DI LOLLO, V., ENNS, J. T., & RENSINK, R. A. (in press). Competition for consciousness among visual events: The psychophysics of reentrant visual pathways. *Journal of Experimental Psychology: General*.
- DUNCAN, J. (1980). The locus of interference in the perception of simultaneous stimuli. *Psychological Review*, *87*, 272-300.
- DUNCAN, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology: General*, *113*, 501-517.
- DUNCAN, J. (1993). Similarity between concurrent visual discriminations: Dimensions and objects. *Perception & Psychophysics*, *54*, 425-430.
- DUNCAN, J., WARD, R., & SHAPIRO, K. L. (1994). Direct measurement of attentional dwell time in human vision. *Nature*, *369*, 313-315.
- ENNS, J. T., BREHAUT, J., & SHORE, D. I. (1996). Attended objects are on view longer than unattended ones. *Abstracts of the Psychonomic Society*, *1*, 40.
- ENNS, J. T., & DI LOLLO, V. (1997). Object substitution: A new form of masking in unattended visual locations. *Psychological Science*, *8*, 135-139.
- GANZ, L. (1975). Temporal factors in visual perception. In E. C. Carterette and M. P. Friedman (Eds.), *Handbook of Perception: Vol. 5. Seeing* (pp. 170-231). New York: Academic Press.
- GIESBRECHT, B. L., & DI LOLLO, V. (1996, November). *Beyond the attentional blink: Visual masking by item substitution*. Poster presented at the annual meeting of the Psychonomic Society, Chicago.
- GIESBRECHT, B. L., & DI LOLLO, V. (1998). Masking in the attentional

- blink: Decoupling integration from interruption masking. *Journal of Experimental Psychology: Human Perception & Performance*, **24**, 1454-1466.
- GRANDISON, T. D., GHIRARDELLI, T. G., & EGETH, H. E. (1997). Beyond similarity: Masking of the target is sufficient to cause the attentional blink. *Perception & Psychophysics*, **59**, 266-274.
- HUSSAIN, M., SHAPIRO, K., MARTIN, J., & KENNARD, C. (1997). Abnormal temporal dynamics of visual attention in spatial neglect patients. *Nature*, **385**, 154-156.
- JOLICOEUR, P., & DELL'ACQUA, R. (1998). The demonstration of short-term consolidation. *Cognitive Psychology*, **36**, 138-202.
- KAHNEMAN, D. (1968). Method, findings, and theory in studies of visual masking. *Psychological Bulletin*, **9**, 404-425.
- LUCK, S. J., VOGEL, E. K., & SHAPIRO, K. L. (1996). Word meanings can be accessed but not reported during the attentional blink. *Nature*, **383**, 616-618.
- MARTIN, J., ISAAK, M. I., & SHAPIRO, K. L. (1995, August). *Probe identification errors support an interference model of the attentional blink in rapid serial visual presentation*. Poster presented at the annual meeting of the American Psychological Society, New York.
- 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.
- POSNER, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, **32A**, 3-25.
- 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 & Performance*, **18**, 849-860.
- RENSINK, R. A., O'REGAN, J. K., & CLARK, J. J. (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science*, **8**, 368-373.
- SCHEERER, E. (1973). Integration, interruption and processing rate in visual backward masking. *Psychologische Forschung*, **36**, 71-93.
- SEIFFERT, A. E., & DI LOLLO, V. (1997). Low-level masking in the attentional blink. *Journal of Experimental Psychology: Human Perception & Performance*, **23**, 1061-1073.
- SHAPIRO, K. L. (1994). The attentional blink: The brain's eyeblink. *Current Directions in Psychology*, **3**, 86-89.
- SHAPIRO, K. L., CALDWELL, J., & SORENSEN, R. E. (1997). Personal names and the attentional blink: A visual "cocktail party" effect. *Journal of Experimental Psychology: Human Perception & Performance*, **23**, 504-514.
- SHAPIRO, K. L., DRIVER, J., WARD, R., & SORENSEN, R. E. (1997). Priming from the attentional blink: A failure to extract visual tokens but not visual types. *Psychological Science*, **8**, 95-100.
- SHAPIRO, K. L., RAYMOND, J. E., & ARNELL, K. M. (1994). Attention to visual pattern information produces the attentional blink in rapid serial visual presentation. *Journal of Experimental Psychology: Human Perception & Performance*, **20**, 357-371.
- SIMONS, D. J. (1996). In sight, out of mind: When object representations fail. *Psychological Science*, **7**, 301-305.
- SPENCER, T. J., & SHUNTICH, R. (1970). Evidence for an interruption theory of backward masking. *Journal of Experimental Psychology*, **85**, 198-203.
- TREISMAN, A. M., & GELADE, G. (1980). A feature integration theory of attention. *Cognitive Psychology*, **12**, 97-136.
- TSAL, Y., MEIRAN, N., & LAMY, D. (1995). Towards a resolution theory of visual attention. *Visual Cognition*, **2**, 313-330.
- TURVEY, M. T. (1973). On peripheral and central processes in vision: Inferences from an information-processing analysis of masking with patterned stimuli. *Psychological Review*, **80**, 1-52.
- WARD, R., DUNCAN, J., & SHAPIRO, K. (1997). Effects of similarity, difficulty, and nontarget presentation on the time course of visual attention. *Perception & Psychophysics*, **59**, 593-600.
- ZUVIC, S. M., VISSER, T. A. W., & DI LOLLO, V. (in press). Direct estimates of processing delays in the attentional blink. *Psychological Research/Psychologische Forschung*.

(Manuscript received June 11, 1997;
revision accepted for publication August 11, 1998.)