Visual Masking and Task Switching in the Attentional Blink

James T. Enns, Troy A. W. Visser, Jun-ichiro Kawahara, & Vincent Di Lollo

University of British Columbia

Running head: Visual Masking and Task Switching

Chapter to appear in Shapiro, K. (Ed.), The limits of attention: Temporal constraints on human information processing. Oxford University Press.

Address correspondence to:
J. T. Enns
Department of Psychology
University of British Columbia
2136 West Mall
Vancouver, BC Canada V6T 1Z4

jenns@psych.ubc.ca
phone: 604-822-6634

Abstract

When two targets are presented in rapid succession, identification of the first target is nearly perfect, but identification of the second target is impaired when it follows the first by less than about 500 ms. In the present chapter, we examine the role of two factors – visual masking and attentional switching – in producing this second-target deficit, commonly known as the attentional blink (AB). We begin by reporting that the AB is affected in different ways by how the first and second targets are masked. The AB occurs whether the first target is masked by interruption or by integration. In contrast, only interruption masking is effective for the second target. In the next section, we examine the role of attentional switching as revealed by the phenomenon of "lag-1 sparing", which refers to a reduction in the magnitude of the AB when the second target is presented directly after the first. This "sparing" points to the importance of attentional switching between sequential tasks in producing the AB. In the final section, we present new evidence showing that when the two targets have different attentional requirements, the AB occurs reliably even without a mask after the second target. Based on this evidence, we conclude that the AB may represent the temporal cost of reconfiguring the visual system in readiness for different sequential tasks.

0. Introduction

This chapter summarizes research, conducted in our laboratory over the past 5 years, on the phenomenon known as the attentional blink (AB). In one popular procedure used to induce the effect, observers attempt to identify two target letters inserted into a temporal stream of digits (Chun & Potter, 1995; Raymond, Shapiro, & Arnell, 1992; Shapiro, 1994). Because all the items are displayed in the same location, one every 100 ms or so, this procedure is often referred to as rapid serial visual presentation (RSVP). The AB refers to a striking imbalance in the identification accuracy of the two targets. Although the first target can be identified almost perfectly, identification of the second target is substantially impaired, especially when it is presented with a temporal lag of 200-500 ms after the first target.

The same pattern of results is obtained using a simplified version of the task (Duncan, Ward, & Shapiro, 1994; Ward, Duncan & Shapiro, 1997) which we call the two-target procedure. In this case, the only items presented are the two targets, displayed in different screen locations at various temporal lags from each other, and each is followed by a pattern mask. As with the RSVP method, identification is nearly perfect for the first target, but is dramatically reduced for the second target when it lags the first by 200-500 ms. The perceptual and cognitive mechanisms responsible for this second-target deficit are the focus of this chapter.

In our research we have come to the conclusion that there are two quite separable aspects to the AB. One aspect concerns the contributions of visual masking. Because the methods for studying the AB typically involve the rapid presentation of visual patterns in the same spatial location, issues of masking cannot be ignored. Indeed, as we will show in summarizing our research, whether or not an AB is observed can depend critically on the kind of masking that is employed.

These studies have also revealed an important relationship between masking and attention. Namely, visual pattern perception is shown to be vulnerable to backward masking only when attention cannot be focused rapidly on a target stimulus. This means that in a typical AB experiment, the second target is inherently more vulnerable to backward masking than the first. By the very design of the task, the first target is presented when the observer is optimally prepared to process it (i.e., it is the first target to engage response selection and execution processes). In contrast, the second target is presented shortly after the first, such that if the observer is not fully prepared to process the second target, its representation will be substituted in low-level visual registers by a representaton of an item that follows the second target. Accuracy in reporting the second target therefore suffers as a direct consequence of what we call an object substitution process for unattended items.

The second aspect of the AB concerns attention, by which we mean the control and deployment of limited capacity mechanisms in the brain involved in the identification of patterns. As we will show, in the absence of any visual masking, whether the AB is observed or not depends critically on the degree to which attentional mechanisms must be reconfigured from the task designated for the first target to the task required for the second target. At one extreme, we will demonstrate that no AB occurs when the tasks for each target have similar attentional demands. At the other extreme, we will show that a large AB ensues, even in the absence of visual masking, if the attentional requirements differ sufficiently between the first and second targets. Because of these differences, we will argue that the AB is best understood, not as a measure of the time needed to perform the first task (i.e., the attentional dwelltime of Duncan et al, 1994), nor as an index of those second-target tasks which can be performed concurrently with the first-target task without attentional conflict (i.e., the concurrent processing measure of Joseph, Chun & Nakayama, 1997). Instead, we believe the AB to be a temporal index of the cognitive cost of reconfiguring the visual processing system for the performance of one task versus another.

We present our story in three sections. In the first section we summarize our already-published research on the multiple roles of visual masking in the AB. In the second section we introduce the basis for our ideas on attentional reconfiguration through an examination of a curious and little understood feature of AB studies: whether or not second targets presented in the temporal position immediately following the first target are subject to the blink. The reduction of AB for the second target in that position is refered to as lag-1 sparing and it can be linked reliably to the differences in processing demands between the first and second target (Visser et al, in press). In the third section, we extend these ideas to the entire range of AB effects across temporal lag. This is the most speculative aspect of our story, since it will depend ultimately on the outcome of a series of experiments that are not yet complete. However, we argue this case by describing selected conditions from ongoing experiments that point to the AB as a measure of the cost of task switching, or attentional reconfiguration.

1. The Role of Masking in the AB

Visual masking refers to a reduction in the visibility of a stimulus, called the target, which occurs as a result of the presentation of a second stimulus, called the mask. Masks can be divided into two broad categories; those consisting of patterns which overlap the targets in space, often referred to as pattern masks, and those that do not (Breitmeyer, 1984). In the present work, we emphasize the role of pattern masking in the AB.

Pattern masks present the visual system with at least two different kinds of spatio-temporal conflict (Breitmeyer, 1984; Ganz, 1975; Kahneman, 1968; Scheerer, 1973; Turvey, 1973). One conflict occurs when the target and mask are perceived as part of a unitary pattern. In this case, masking occurs by a process of integration and is akin to the addition of noise (the mask) to the signal (the target). Turvey (1973) referred to this as "peripheral masking" to emphasize that it takes place at early levels of visual processing. A second kind of conflict arises when processing of a first pattern (the target) is interrupted by a second pattern (the mask) which appears in the same spatial location before the target has been fully processed. This conflict involves a competition for the higher-level mechanisms required for object recognition. Therefore, it is often referred to as masking by interruption or "central masking" (Turvey, 1973). Unlike integration masking, interruption masking occurs only when the mask follows the target in time, with target accuracy lowest at target-mask intervals that are greater than zero, and gradual improvement as interval increases (e.g., Bachman & Allik, 1976; Turvey, 1973).

1A. Masking of the First Target: Any Mask Will Do

Masking of the first target emerged as an important factor in early demonstrations of the AB (Raymond, Shapiro, & Arnell, 1992). These investigators found that the AB did not occur when the item directly following the first target (the "+1 item") was replaced by a blank screen. On the hypothesis that the +1 item acted as a mask, its omission allowed processing of the first target to continue without interference for the ensuing mask-free period.

In an initial experiment aimed at demystifying the role of first-target masking, Seiffert and Di Lollo (1997) examined the effects of integration masking by presenting the first target and the +1 item simultaneously and spatially superimposed. This combined stimulus was followed by a blank screen during the interval in which the +1 item would normally have appeared. A large AB was obtained as compared to a condition in which the first target was not masked (i.e. the +1 item was omitted), suggesting that integration masking of the first target was one way in which an AB could be obtained.

Having established a role for integration masking, Seiffert and Di Lollo (1997) went on to evaluate another form of masking. They did this by presenting the +1 item directly after the first target, but in a location that was spatially adjacent to where the target had appeared. This arrangement eliminated the possibility of masking by integration, but was suitable for metacontrast masking, in which processing of the target is interrupted by a trailing non-overlapping mask. Again, the magnitude of the AB was significantly greater than that produced when the first target was not masked, a result also found by Grandison, Ghirardelli, and Egeth (1997). Collectively, these results led Seiffert and Di Lollo to conclude that, in order to obtain an AB, the first target must be masked. The precise form of masking, however, is unimportant. An AB occurs whether the first target is masked by integration, interruption, or even metacontrast.

Masking of the first target is clearly important in the production of the AB. However, it is not immediately obvious why it should affect identification of the second target. Indeed, this is quite surprising, given that masking of the first target impairs identification of the second target more than identification of the first. The explanation proposed by both Seiffert and Di Lollo (1997) and Grandison et al. (1997) follows from the two-stage theory of the AB proposed by Chun and Potter (1995). This theory comprises an initial stage in which items are detected rapidly as potential targets, and a second capacity-limited stage in which these items are processed in greater detail for subsequent report. Access to Stage 2 is gained by items that have been identified as potential targets in Stage 1. Until processing of the first target is completed, however, no subsequent items can gain access to Stage 2. Therefore, if the second target is presented before Stage 2 is free, access to Stage 2 will be delayed. The AB occurs because the initial representation of second target decays or is overwritten by subsequent items during the delay. On this account, masking increases the length of time required to process the first target. Although this does not ultimately impair first-target accuracy, it does delay the admission of the second target into Stage 2. As the delay increases, the likelihood that the second target will decay or be overwritten in Stage 1 also increases (Grandison et al., 1997; Seiffert & Di Lollo, 1997).

This account of the influence of first-target masking is consistent with the general relationship between difficulty of processing the first target and the AB posited earlier by Chun and Potter (1995). From this perspective, masking is simply one of many potential ways to make processing of the first target more difficult, and to increase the period of delay for the second target. One implication of this account is that the size of the AB should vary inversely with first-target accuracy. To evaluate this relationship, Seiffert and Di Lollo (1997) compared the percentage of correct responses on the first target with the magnitude of the AB using their own data as well as that from 26 other separate experiments. The results of this analysis are presented as a scatter diagram in Figure 1. As predicted, there was a significant negative correlation between first target accuracy and the magnitude of the AB , $r = 0.73$, $p < .001$.

---------------- Insert Figure 1 about here ----------------

The proposed relationship between difficulty of processing the first target and the magnitude of the AB has not gone unchallenged. Ward et al (1997) noted that all of the studies cited by Seiffert and Di Lollo (1997) manipulated difficulty through masking. This left open the possibility that other difficulty manipulations might not produce the same relationship. To test this, Ward et al. (1997) varied the difficulty of a size judgement for the first target. They found that although accuracy was reduced when the judgement task was made more difficult, the magnitude of the AB was unaffected. On the basis of this result, Ward et al. (1997) concluded that not all manipulations of first target difficulty lead to changes in the AB, and that earlier findings might be specific to masking.

1B. Masking of the Second Target: Interruption is Essential

It has long been recognized that the second target must be masked in order for the AB to occur. Ostensibly, the purpose of masking has been to increase the difficulty of processing the second target, thereby bringing accuracy within a measurable range (e.g. Moore, Egeth, Berglan, & Luck, 1996). If this were the principal function served by masking, then either integration or interruption masking should be sufficient, and an AB should be found using either procedure.

To investigate the the two forms of masking, Giesbrecht and Di Lollo (1998) compared the magnitude of the AB when the second target was masked by interruption or by integration. Both targets were letters, presented in an RSVP stream of digit distractors. In the interruption-masking condition, the second target was followed by at least one digit that acted as a mask. In the integration-masking condition, the second target was presented simultaneously with a digit, with no additional items following in the RSVP stream. This brought accuracy of the second target within a measureable range, while eliminating interruption masking by trailing items.

When the second target was masked by interruption, a significant AB was obtained, with identification of the second target steadily improving beyond an inter-target interval of about 200 ms. This finding was consistent with results obtained in other studies (e.g. Chun & Potter, 1995; Raymond et al., 1992; Shapiro, Raymond, & Arnell, 1994). In contrast, when the second target was masked by integration, identification was impaired equally across all lags, but the lag-dependent deficit that is the signature of the AB was missing.

Similar results were obtained by Brehaut et al. (in press) using the two-target procedure. As shown in Figure 2, when the second target and the mask were presented simultaneously, a lag-dependent deficit was notably absent. However, as the interval between target and mask increased, an AB became increasingly apparent. On the basis of these findings, it is clear that interruption masking of the second target is more than a methodological convenience – it is necessary to obtain an AB .

----------------- Insert Figure 2 about here -----------------

We have argued that the two-stage theory (Chun & Potter, 1995) provides a comprehensive account for these results (Brehaut et al, in press; Giesbrecht & Di Lollo, 1998). As noted earlier, this theory specifies two sources of deterioration for the second target while it is delayed in Stage 1: passive decay and overwriting by temporally trailing items. The results of Giesbrecht and Di Lollo (1998) suggest that the main source of deterioration in Stage 1 is overwriting. Passive decay would be evidenced by impairment in accuracy at the shortest lag, followed by a gradual improvement over lags. If decay was an important determinant of accuracy, this trend should have been observed in the integration masking condition in which there were no trailing items to erase the second target. However, the results in this condition show little or no evidence of such a trend. This suggests that passive decay does not play a major role in the deterioration of the second target.

In contrast, overwriting provides a consistent account of the results obtained with both integration and interruption masking. With integration masking, accuracy was impaired because the pattern of the second target was impoverished, making it harder to extract the target from the noise. However, the magnitude of the impairment did not vary as a function of lag. This result can be explained within the two-stage theory by noting that the second target and the mask formed a unitary stimulus that remained available in Stage 1 because there was no trailing item to overwrite it. Thus, when the delay experienced by the second target ended, its representation entered Stage 2 and was processed to the extent allowable by its reduced stimulus quality. On the other hand, with interruption masking, second target accuracy did vary with lag. Accuracy was most impaired at lags of 200-300 ms, and improved progressively thereafter. According to the two-stage theory, accuracy was impaired at the shorter lags because the trailing mask erased the second target during the period of delay in Stage 1. At longer lags, the probability that the second target could enter Stage 2 before being erased by the mask increased, and accuracy improved accordingly.

Implicit within this account is an object substitution theory of interruption masking. While it is delayed in Stage 1, the representation of the second target is vulnerable to overwriting by the trailing mask. When that happens, the representation of the mask replaces that of the second target and eventually gains access to Stage 2. As a consequence of this replacement, the mask is substituted for the second target as the object for eventual conscious registration. This object substitution theory argues that the interruption mask does more than simply halt processing of the second target. Instead, because the representation of the target has been overwritten, Stage 2 mechanisms are left only with the mask to process.

A view of interruption masking akin to object substitution is well supported in the masking literature. There is ample evidence to show that when two targets are presented sequentially at an optimal interval, it is the second that is perceived to the detriment of the first (Bachmann & Allik, 1976). This effect has been found to be more pronounced in unattended visual locations, suggesting that stimuli displayed outside the focus of attention are more likely to be delayed in Stage 1, thus remaining vulnerable to substitution over a longer period (Enns & Di Lollo, 1997).

This tendency has been observed not only when attention is distributed over space, but also when it is distributed over time, as in the AB . In studies by

Chun (1997) and Martin, Issak, and Shapiro (1995), an RSVP stream of letters was presented with the two target letters being distinguished from the rest of the stream by either colour (Chun, 1997) or size (Martin et al., 1997). In both studies, the most prominent type of error consisted of misidentifying the second target as the item directly following it. These results are consistent with the hypothesis that while delayed in Stage 2, the second target is overwritten by the trailing item.

1C. Summary of Masking in the AB

Considered collectively, the results demonstrate that masking plays two separate roles in the AB. Masking of the first target, whether by integration, interruption, or metacontrast, introduces a delay in the processing of the second target. During this delay, the second target is vulnerable to being replaced by a trailing mask with a probability that decreases as the temporal lag between the targets is increased. These findings suggest that the function of first-target masking is to make its processing more difficult and that several kinds of masking accomplish this equally well.

The second role played by visual masking is specific to the second target. Here, it is essential to use interruption masking in order to obtain the AB. Masking by integration, although sufficient to reduce second target accuracy, does not produce a lag-dependent effect. In addition to being of methodological interest, this finding points strongly to a visual process of considerable theoretical importance. This process, referred to as object substitution, is the overwriting or replacement of visual representations by trailing items, when those representations are unattended. Given how much of our visual world goes unattended, object substitution may play an extensive role in perception.

Indeed, we have suggested (Brehaut et al., in press) that substitution may underlie demonstrations of "change blindness," referring to the finding that major changes in the visual world can go undetected if attention is misdirected away from the area of change (Rensink, O'Regan, & Clark, 1997; Simons, 1996).

2. Lag-1 Sparing and the Attentional Blink

The AB points to the limits in the ability of the visual system to process information. When two targets are presented in rapid succession, processing resources required in common by both targets are less available for the second target until processing of the first has been completed. On this account, identification of the second target should be maximally impaired when it is presented directly after the first target, in the ordinal display position known as Lag 1. The deficit should then diminish as lag increases, reflecting the increased availability of resources previously deployed to the first target. However, this pattern of results is found in only about one-half of published AB experiments. In the other half, accuracy is largely unimpaired at lag 1, drops substantially at lags 2 and 3, and then gradually recovers. It is this pattern of improved accuracy at lag 1, followed by a pronounced deficit at longer lags, that is referred to as lag-1 sparing (Potter, Chun, Banks, & Muckenhoupt, 1998).

Lag-1 sparing has been attributed to the sluggish closing of an attentional gate (Chun & Potter, 1995; Shapiro & Raymond, 1994). The gate is said to open upon the presentation of the first target and then to close slowly, thus allowing a trailing stimulus to gain access to processing resources along with the first target. If the trailing item is the second target, then it is processed along with the first target and an AB is avoided. Although this explanation can account for the occurrence of lag-1 sparing, it cannot explain why lag-1 sparing does not occur, under identical rates of presentation, in about one-half of studies. These failures to obtain lag-1 sparing suggest that factors other than temporal contiguity must determine whether sequential stimuli give rise to the same or different attentional episodes. By uncovering these factors, the phenomenon of lag-1 sparing may become relevant to broader issues relating to distribution of attention and rapid changes in attentional set in both spatial and non-spatial domains.

In studying the dynamics of switching attentional sets, the main issue is how the cognitive system is reconfigured to cope with rapidly changing demands. The impact of such reconfiguration can be observed readily in studies of the AB, because observers must perform a rapid attentional switch from the first target to the second across a brief temporal gap. In particular, accuracy in the report of the second target at lag 1 indicates whether the intervening interval of approximately 100 ms is sufficient for an attentional switch to be made. Lag-1 sparing suggests that a switch was successful, with little detriment to second target accuracy. Absence of lag-1 sparing, on the other hand, suggests that a switch could not be made, perhaps because it required too great a reconfiguration of the system. On this reasoning, identifying the factors that influence lag-1 sparing should reveal factors that are relevant to switching of attentional sets.

Visser, Bischof, and Di Lollo (in press) conducted a systematic examination of lag-1 sparing in the AB literature. Finding over 100 separate experiments in which an AB had been reported, they devised a taxonomic scheme that specified four major classes of switches in attentional set between targets. These included switches in categorical identity (e.g. digits vs letters), in task (e.g. detection vs. identification), in modality (e.g. auditory vs. visual), and in spatial location (e.g. central vs. peripherally-located targets). Visser et al. (in press) then tabulated the frequency of lag-1 sparing as a function of the number and types of switches in attentional set.

The results of this survey are presented in Figure 3. The major trends can be characterized as follows. Lag-1 sparing occurred reliably when there were no switches between targets, or when the switch was unidimensional, involving either task or category alone. Lag-1 sparing was not found with switches in location or with concurrent switches involving two or more categories (e.g. switches of both task and category). As noted by Visser et al. (in press), the presence of multiple switches seemed to act synergistically to prevent lag-1 sparing. For example, task and category switches alone yielded lag-1 sparing in 76% of cases, but when these switches were implemented concurrently, lag-1 sparing occurred in only 18% of cases. These results point to an important relationship between lag-1 sparing and attentional switches that are implemented between targets. At one extreme, lag-1 sparing occurs reliably when there is no switch in attentional set. At the other extreme, lag-1 sparing never occurs when there are multiple switches in attentional set.

---------------- Insert Figure 3 about here ------------------

This pattern of results can be explained by a revised version of the two-stage theory (Chun & Potter, 1995). As noted earlier, this theory suggests that lag-1 sparing occurs when the first and second targets both fit through the same attentional gate. This gate opens upon the onset of the first target, and closes 150-200 ms later. If the second target enters the gate, both targets will become part of the same attentional episode, and both will gain access to high-level mechanisms required for stimulus identification and response planning. As indicated by the large number of studies that have failed to find lag-1 sparing, however, temporal contiguity between targets cannot be the sole criterion for the second target to enter the same attentional gate. Rather the findings reviewed by Visser et al. suggest that an additional criterion is necessary in the form of a filter that controls access to the attentional gate.

The characteristics of this filter are indicated by the relationship between lag-1 sparing and the type of attentional switch implemented between the two targets. To enter the same attentional gate as the first target, the second target must arrive no more than 150-200 ms later and it must match the characteristics of the input filter. If both of these conditions are met, the two targets will become part of the same attentional episode and will gain access to high-level processing mechanisms. In this event, lag-1 sparing will occur. If the second target does not pass the input filter, perhaps because it differs from first target on some dimensions, an appropriate new filter needs to be set up. By the time this new filter is ready, the temporal gate will have closed, and the second target will remain in Stage 1 where it is vulnerable to masking. In this event, lag-1 sparing does not occur because the representation of the second target will have been replaced by the that of the mask by the time Stage 2 is again free.

To function optimally, the input filter must operate at a relatively low level in the visual system, where sensory signals flowing towards higher cortical centres can be monitored. At the same time, the filters must be responsive to rapid changes in attentional set and in response planning, which are functions normally associated with high-level structures in the prefrontal cortex (Goldman-Rakic, 1987, 1988). These requirements suggest that the gating mechanism may be instantiated in more than one brain area, depending on the nature of the attentional set. For example, attending selectively to stimuli in motion would probably include gating circuitry in cortical areas V1 and V5. But, there are strong indications that input filters can be set up even more peripherally than primary visual cortex. Among the likely candidates for subcortical gating mechanisms is the perigeniculate nucleus, which is a network interposed between the lateral geniculate nucleus and area V1. Sitting astride the main input pathways to the visual cortex, the perigeniculate nucleus is ideally suited for monitoring incoming sensory signals.

What is suggested by these considerations is an intelligent filtering system that is not restricted to solely physical features. Rather the reconfiguring of a filter is probably part of a more comprehensive and goal-oriented process aimed at selecting those stimulus attributes that are useful for performing the task at hand. Such task-set reconfigurations have often been associated with higher cortical functioning (Monsell, 1996). However, in view of the reciprocal neural connectivity between prefrontal cortex and other brain regions such as the perigeniculate nucleus, intelligent filtering, even at the earliest input level, seems to be a realistic possibility (Motter, 1993; Roelfsema, Lamme, & Spekreijse, 1998; Somers, Dale, Seiffert, & Tootell, 1999).

The idea of intelligent filtering invites a reversal in the usual conceptualization of events in visual information processing. Traditionally, such processing has been thought to occur in two discrete stages (e.g., Neisser, 1967). The first is an input stage of nearly unlimited capacity in which stimuli are encoded in parallel. This stage is considered to be largely stimulus-driven, and hence under exogenous control. The second is a resource-limited stage that processes stimuli serially under the direction of a unitary controlling mechanism. In contrast to the first stage, processing at this second stage is considered to be conceptually-driven, and hence under endogenous control.

Contrary to this viewpoint, the intelligent filtering proposed by Visser et al. suggests that the loci of exogenous and endogenous control be reversed. The input filters are reconfigured under the control of signals from brain areas as high as prefrontal cortex, and their configuration determines whether any given stimulus can gain access to higher processing levels, or whether it is excluded. Given the role played by prefrontal cortex, it follows that the functioning of this input stage must be governed largely by endogenous, conceptually-driven signals.

Conversely, it would seem that the second-stage is governed indirectly by exogenous stimulus-driven events. Consider that if an item passes the input filters, it gains access to this second stage which consists of specialized feature- or task-specific processing modules. What processing module is activated by this item will depend on the nature of the stimulus. For example, stimuli in motion will activate motion-processing modules, but may not activate colour-processing modules. It follows from these considerations that the functioning of processing modules in this second, high-level stage must be governed in good part by exogenous, stimulus-bound factors.

The functional organization of the brain implied by this viewpoint is that of a number of independent, inter-connected processors, each of which performs a specific function. These independent processors operate in parallel on incoming stimuli that have passed the input filters. This is consistent with view expressed by Allport, Antonis, and Reynolds (1992), Allport, Styles, & Hsieh

(1994), and Monsell (1996), and is congruent with the modular brain organization revealed in current neuroanatomical and neurophysiological studies (e.g. Felleman & Van Essen, 1991; Posner & Raichle, 1994).

3. Task Switching and the Attentional Blink

From the foregoing discussion, two factors have emerged as important in the AB: backward masking of the second target, and attentional switching between targets. What is less clear from the extant literature is how these two factors combine in producing an AB. One way to examine their relationship is illustrated in Figure 4, where the two factors are combined in a 2x2 table.

---------------- Insert Figure 4 about here -----------------

Readily noticeable in Figure 4 is the absence of studies in which an attentional switch was implemented, but the second target was not masked. This lacuna is understandable because, in most studies of the AB, the second target has been masked for a practical reason: unless it was masked, accuracy was at ceiling, and an AB was no longer in evidence. An exception to this rule was seen when the second target, instead of being backward-masked, was degraded by visual noise (Brehaut et al., in press; Giesbrecht and Di Lollo, 1998). Under these conditions, a lag-dependent deficit in second-target accuracy was not obtained, leading to the conclusion that backward masking of the second target was necessary for obtaining an AB.

This conclusion is consistent with the evidence from all other AB studies, but its generality is constrained in that it was based on a study in which there was no attentional switch between targets. Specifically, in the study by Giesbrecht and Di Lollo (1998), both targets were alphabetical characters presented in the same spatial location. For this reason, the conclusion reached in that study cannot be generalized to cases in which an attentional switch is required. What remains to be determined, then, is whether an AB occurs when an attentional switch is implemented between the two targets, but the second target is not followed by a masking stimulus.

A study designed for that purpose was carried out by Kawahara, Enns, and Di Lollo (1999). Its major objective was to determine whether an attentional switch between targets is sufficient to produce an AB , even if the second target is not masked. Observers viewed an RSVP stream of digit distractors displayed in the centre of the screen, and were required to identify both targets. The first target was always a letter, and the second was either another letter (no switch) or a diagonal line whose direction of tilt was to be identified (attribute switch). In addition, the second target was displayed either in the same location as the first (no switch), or unpredictably in one of 12 eccentric locations (location switch). To bring the level of second-target identification within a measurable range, the second target was embedded in random-dot noise, but was never followed by a mask. The design was a 2x2 between-subject factorial, in which the presence or absence of a switch in location was crossed with the presence or absence of a switch in attribute.

---------------- Insert Figure 5 about here ----------------

The results are shown in Figure 5, averaged over 16 observers in each of the four groups. In agreement with the findings of Giesbrecht and Di Lollo (1998), no statistically significant AB was obtained when there was no switch between targets (Figure 5A). In contrast, a significant AB was obtained in all conditions involving attentional switches, despite the fact that the second target was never followed by a mask. In addition, the AB in Figure 5D was reliably larger than that in Figure 5B, suggesting that a concurrent switch in two dimensions (attribute and location) produced a larger AB than a switch in location alone. A compelling inference from these results is that an attentional switch between targets is sufficient for producing an AB , even when the second target is not followed by a mask. This information can be entered in the empty cell in Figure 4, thus completing the 2x2 table. It is now apparent from Figure 4 that an AB is obtained with any combination of second-target masking and attentional switching, but is not obtained when there is neither masking nor switching. We next consider the relative roles of attentional switching and masking in producing the AB .

3A. On the role of attentional switching

It is clear from Figure 4 that an attentional switch between targets leads to an AB even when the second target is not masked. Why is perception of the second target impaired by such a switch? Our answer to this question hinges on the assumption that an attentional switch entails a reconfiguration of the visual system. The idea of a system reconfiguration was outlined earlier in this chapter, and has been discussed extensively by Monsell (1996) and by Visser et al. (in press). In brief, what is reconfigured is an intelligent filtering system which operates on the visual input to select and encode task-relevant items, and to distinguish them from other incoming stimuli.

When the task requires an attentional switch between targets, the configuration of the system must be changed from one tuned to the characteristics of the first target to one tuned to those of second. If the second target arrives before the system has been suitably reconfigured, an adequately-encoded representation of the second target cannot be formed, and its identification is correspondingly impaired. Thus, an AB occurs even if the second target is not masked, because an attentional switch prevents a representation of it from being formed at the earliest stage. On this option, the AB shown in the lower two panels in Figure 4 would be mainly attributable to the processing deficit arising from an attentional switch. Under these conditions, the presence or absence of a trailing mask would be largely irrelevant. Our reasoning is that if a representation of the second target was not formed to begin with, there would be nothing of consequence to be overwritten by a trailing mask.

This proposal of attentional reconfiguration leads to an interesting prediction. When an observer views an RSVP stream and is required to make an attentional switch between the first and second target, there should be no priming of a third target by the second target, as has been found by Shapiro, Driver, Ward, & Sorensen (1997) in a task involving no switch between the first and second targets. When there are no switches of this kind, the mechanisms activated by the second target prior to being masked can serve to prime the processing of the third target. However, when there is a switch, a representation of the second target will not even be formed. Accordingly, the mechanisms for the second target are never activated and so no priming would be expected.

An alternative interpretation to that of attentional switching should be considered and dismissed. It may be suggested that an adequate representation of the second target is formed at an early processing stage even when the task requires an attentional switch. During the time period of the switch, however, that representation undergoes rapid decay and becomes undecipherable by the time the system has been suitably reconfigured. The ensuing AB would then be mediated by the decay of the second target. This option is disconfirmed by two findings in the literature. Reaction time studies have shown that processing of the second target is substantially delayed while the system is processing the first target (Arnell & Duncan, 1997; Jolicoeur, 1998). Were the representation of the second target to decay during this delay, a lag-dependent deficit should be expected. Yet, no such deficit has been revealed in studies in which the second target was not masked, and level of second-target identification was well below ceiling (Figure 5a; Brehaut et al, in press; Giesbrecht & Di Lollo, 1998). This all but rules out decay as a factor in the AB.

3B. On the role of second-target masking

It is clear from Figure 4 that backward masking of the second target invariably yields an AB, whether or not the task involves an attentional switch between targets. Despite the consistency of these results, a complete account of the role of masking is not entirely straight-forward. When there are no attentional switches (Figure 4, upper panels), the evidence is unambiguous: an AB is obtained only if the second target is followed by a mask. We have seen how this can be explained by the joint actions of backward masking and inattention. Namely, while the first target is being processed, the second target remains unattended and vulnerable to overwriting by the trailing mask.

On the other hand, when there is an attentional switch, an AB is invariably obtained, even if the second target is not followed by a mask (Figure 4, lower panels). Clearly, masking cannot be regarded as necessary under these conditions. But does it still play a role in bringing about an AB? This question cannot be answered on the basis of extant information. There are at least two options. First, the mask may destroy the early representation of the second target, thus rendering the consequences on an attentional switch irrelevant. The second option is that an attentional switch entails a reconfiguration of the system such that the second target fails to even be encoded. This, by itself, would produce a lag-dependent deficit, and would render a trailing mask irrelevant. On this option, we would expect that either integration or interruption masking of the second target would produce an AB, provided that there was an attentional switch. This would contrast with the our earlier finding (Brehaut et al, in press; Giesbrecht & Di Lollo, 1998) that interruption masking of the second target is necessary to obtain an AB when there is no attentional switch.

4. Concluding Remarks

The view that has come to guide our research on the attentional blink is one in which “attention” refers to the coordinated effort of a widely distributed group of specialized brain regions. These include lower level registers of exclusively visual information (LGN and V1), higher level centers involved in pattern analysis and visual object identification (visual temporal cortex) and supramodal centers involved in the selection, coordination and prioritization of information flow leading to action (thalamus, prefrontal cortex). Described this way, the approach bears a close resemblance to that of several other widely-held views of attention (Colby, 1991; Desimone & Duncan, 1995; Duncan, 1996; Posner & Raichle, 1994).

There are two features of our proposal, however, that we believe are unique. The first is a deliberate attempt to incorporate modern anatomical and physiological data concerning reciprocal and reentrant connections between specialized brain regions into our accounts and predictions regarding behavior. For example, the idea of intelligent filtering proposed by Visser et al. (in press) is premised on the functional importance of reentrant neural pathways from cortical centers of attentional control (e.g., frontal cortex, parietal cortex) to centers of attentional expression (e.g., LGN, area V1). In other work, we have described how we believe reentrant neural connections are responsible for object substitution masking effects, both in the attentional blink (Geisbrecht and Di Lollo, 1998) and in metacontrast masking (Di Lollo, Enns & Rensink, under review).

The second unique aspect of our proposal concerns its implications for what is being measured in an AB experiment. We have tried to make it clear that the AB, observed in the absence of any masking of the second target (the right side in Figure 4), is an index of the cost incurred by the visual system in reconfiguring itself for a task (the second target) that is different from a recently completed task (the first target). This means that the AB cannot, even in principle, be used to index the time required to perform the task of identifying the first target, as espoused by the attentional dwelltime view of Duncan et al, 1994). The view of AB as a measure of dwelltime is premised on the assumption that the task of identifying the first target is indexed by its negative influence on the processing of the later-arriving second target. But as we have seen, whether or not an AB is observed for a given first target is dependent on whether an attentional switch must be made to process the second target. The dwelltime view would therefore force the untenable conclusion that the time required to process the first target is dependent on the processing of the later arriving target. Our studies of the AB under conditions of no masking (Figure 5) suggest it is much more likely that the AB is an index of the cost of reconfiguring brain mechanisms to process the second target.

The reconfiguration hypothesis is also at odds with the AB being used as an index of second-target tasks that can be performed without attention (e.g., Joseph, Chun, & Nakayama, 1997). This use of the AB is premised on the assumption that the processing of the second target is uninfluenced by the nature of the processing required for the first target. The question is simply whether the second target task can be performed concurrently with an attention-demanding first target task or not. But this assumption also flies in the face of the dependency of the AB on the relationship between the first and second target tasks. The second target task may, as we have shown, either result in an AB or not, depending on whether a switch in target location is involved (upper two panels in Figure 5) or a switch in target attribute is involved (left side of Figure 5). Again, the data are more consistent with the AB as a measure of system reconfiguration than as a measure of the attentional requirements of the second target.

It is also worth noting, at least in passing, that our emerging view of the AB as a measure of system reconfiguration has interesting implications regarding ecological validity. Consider that almost all the AB research to date has required very little by way of reconfiguration for the observer to process each of the two targets. The large bulk of the studies have used target pairs consisting of two letters or two digits in the same spatial location. This can be contrasted with the typical sequence of tasks we encounter in everyday use of the attentional system. Many times during the day we are required to perform serial target identification tasks such as glancing from the road to the speedometer while driving, glancing from our notes to a classroom of students while teaching, and alternating between visual attention to a ball and opponents in athletic competitions.

Considerations such as these encourage us to study the temporal dynamics of attention under a much wider range of task switching conditions than has been the norm.

We readily acknowledge that the issues we have outlined in our discussion of the AB will be decided ultimately by subjecting them to the strong light of further empirical research. Our primary goal in this chapter has been to make our assumptions, intuitions, and research strategies as explicit as possible, both to ourselves and others, with the hope of making the empirical search for understanding an efficient one. Regardless of the eventual status of various theories concerning the attentional blink, it is already clear that enormous value has been gained in our practical understanding of the temporal dynamics of visual attention.

References

Allport, A., Antonis, B., & Reynolds, P. (1972). On the division of attention: A disproof of the single channel hypothesis. Quarterly Journal of Experimental Psychology, 24, 225-235.

Allport, A., Styles, E. A., & Hsieh, S. (1994). Shifting intentional set: Exploring the dynamics of tasks. In C. Umilt‡ & M. Moscovitch (Eds.), Attention and Performance XV, (pp. 421-452). Hillsdale, N.J.:Erlbaum.

Arnell, K., & Duncan, J. (1997, November). Speeded responses or masking can produce an attentional blink. Poster session presented at the 38th Annual Meeting of the Psychonomic Society, Philadelphia, PA.

Bachmann, T., & Allik, J. (1976). Integration and interruption in the masking of form by form. Perception, 5, 79-97.

Brehaut, J. C., Enns, J. T., & Di Lollo, V. (in press). Masking plays two roles in the attentional blink. Perception & Psychophysics.

Breitmeyer, B. G. (1984). Visual masking: An integrative approach. New York: Oxford University Press.

Chun, M. M. (1997). Temporal binding errors are redistributed in 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 and Performance, 21, 109-127.

Colby, C. L. (1991). The neuroanatomy and neurophysiology of attention. Journal of Child Neurology, 6, S90-S118.

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

Di Lollo, V., Enns, J. T., & Rensink, R. A. (under review). Competition for consciousness among visual events: The psychophsyics of reentrant pathways.

Duncan, J. (1996). Cooperating brain systems in selective perception and action. In T. Inui & J. McClelland (Eds.), Attention & Performance, Vol. XVI (pp. 549-578). Cambridge, MA: MIT Press.

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

Enns, J.T., & Di Lollo, V. (1997). Object substitution: A new form of visual masking in unattended visual locations. Psychological Science, 8, 135-139.

Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in primate visual cortex. Cerebral Cortex, 1, 1-47.

Ganz, L. (1975). Temporal factors in visual perception. In E. C. Carterette & M. P. Friedman (Eds.) Handbook of Perception, Vol. 5, (pp. 169-231). New York: Academic Press.

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

Goldman-Rakic, P. S. (1987). Circuitry of primate prefrontal cortex and regulation of behavior by representational knowledge. In F. Plum & V. B. Mountcastle (Eds.), Handbook of physiology, Section 1, The nervous system, Vol. 5, Higher functions of the brain (pp. 373-417). Bethesda, MD: American Physiological Society.

Goldman-Rakic, P. S. (1988). Changing concepts of cortical connectivity: Parallel distributed cortical networks. In P. Rakic & W. Singer (Eds.), Neurobiology of the neocortex (pp. 177-202). Berlin: Wiley.

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.

Jolicoeur, P. (1998). Modulation of the attentional blink by on-line response selection: Evidence from speeded and unspeeded $Task_1$ decisions. Memory & Cognition, (in press).

Joseph, J. S., Chun, M.M., & Nakayama, K. (1997). Attentional requirements in a preattentive feature search task. Nature, 387, 805-807.

Kahneman, D. (1968). Method, findings, and theory in studies of visual masking. Psychological Bulletin, 70, 404-425.

Kawahara, J., Enns, J. T., & Di Lollo, V. (in preparation). Location and attribute switching in the attentional blink.

Martin, J., Isaak, M. I., & Shapiro, K. L. (1995). 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, NY.

Monsell, S. (1996). Control of mental processes. In V. Bruce (Ed.), Unsolved mysteries of the mind: Tutorial essays in cognition (pp. 93-148). Howe, Sussex: Erlbaum (UK) Taylor & Francis.

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.

Motter, B. C. (1993). Focal attention produces spatially selective processing in visual cortical areas V1, V2, and V4 in the presence of competing stimuli. Journal of Neurophysiology, 70, 909-919.

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

Posner, M. I., & Raichle, M. E. (1994). Images of mind. NY: Scientific American Library.

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

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.

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.

Roelfsema, P. R., Lamme, V. A. F., & Spekreijse, H. (1998). Object-based attention in the primary visual cortex of the macaque monkey. Nature, 395, 376-381.

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 and Performance, 23, 1061-1073.

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 RSVP. Journal of Experimental Psychology: Human Perception and Performance, 20, 357-371.

Simons, D.J. (1996). In sight, out of mind: When object representations fail. Psychological Science, 7, 301-305.

Somers, D. C., Dale, A. M., Seiffert, A. E., & Tootell, R. B. H. (1999). Functional MRI reveals spatially specific attentional modulation in human primary visual cortex. Proceedings of the National Academy of Science, 96, 1663-1668.

Turvey, M. T. (1973). On peripheral and central processes in vision: Inferences from an information-processing analysis of masking with patterned stimuli. Psychological Review, 81, 1-52.

Visser, T. A. W., Bischof, W. F., & Di Lollo (in press). Attentional switching in spatial and non-spatial domains: Evidence from the attentional blink. Psychological Bulletin.

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.

Author Notes

The preparation of this chapter, and the research reported in it, were sponsored by grants from the Natural Science and Engineering Research Council of Canada to the first and last authors.

Figure Captions

Figure 1. Scatter diagram relating attentional blink (AB) magnitude to percentage of corrent responses on the first target. AB magnitude increases significantly with a reduction in first target (T1) accuracy (From Seiffert & Di Lollo, 1997).

Figure 2. Mean accuracy of second target identification under three target-mask intervals (0 ms, 60 ms, 90 ms). Lag-dependent effects increased as the interval between target and mask increased. Chance accuracy level in this task was 5% (Redrawn from Brehaut et al, in press).

Figure 3. Frequency of lag-1 sparing and no lag-1 sparing in the AB literature, plotted as a function of the type of task switch required between the first and second targets. The likelihood of lag-1 sparing decreases directly as a function of task switching (Drawn based on tabular data reported in Visser et al, in press).

Figure 4. A 2x2 table illustrating the results of AB experiments in which the influences of second-target masking and task switching were combined. The results for the experiment identified in the lower right cell are shown in Figure 5.

Figure 5. Mean accuracy of second target (T2) identification when location switching (no, yes) and target attribute switching (no, yes) were combined orthogonally. No AB was observed in the no attribute-no location switching condition (upper left cell). From Kawahara, Enns & Di Lollo (in preparation).

Figure 2

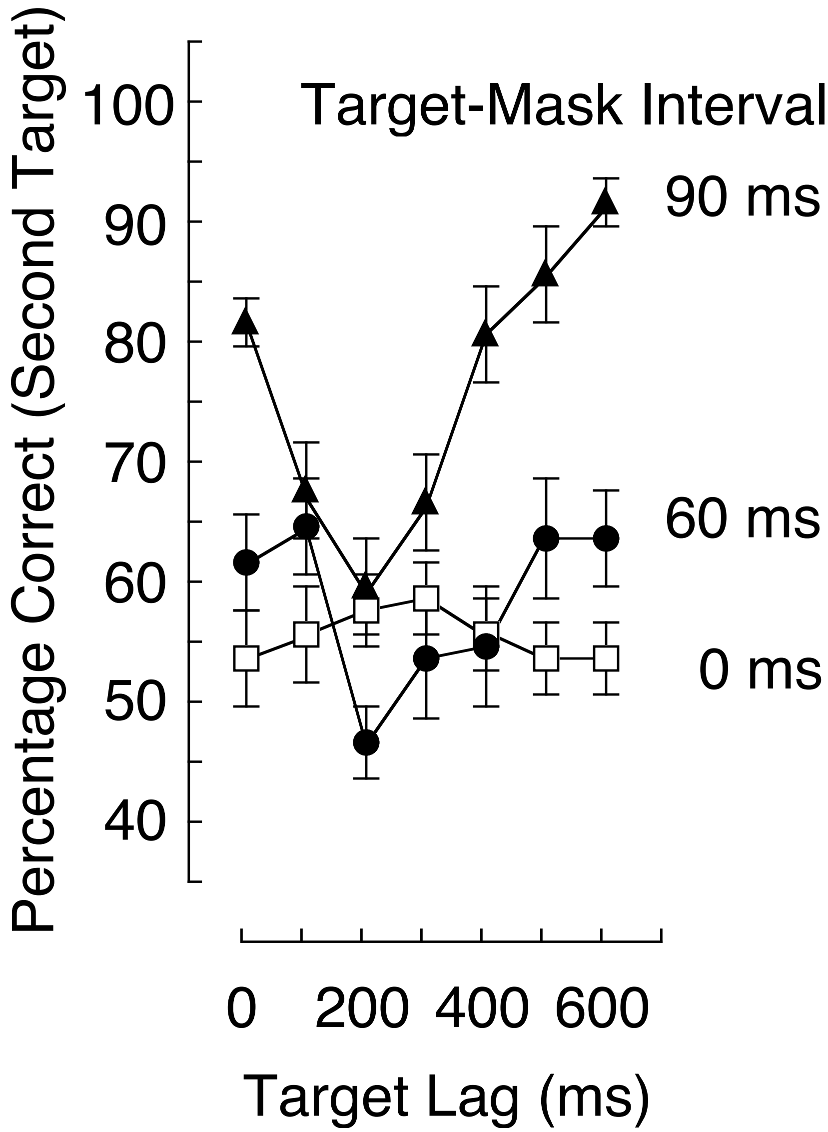
Percentage Correct (Second Target)
100
90
80
70
60
50
40
Target-Mask Interval
90 ms
60 ms
0 ms
0
200
400
600
Target Lag (ms)