# TEMPORAL PERCEPTION IN VISION:

# AN EXAMINATION OF BOTTLENECK MODELS

by

# TROY ANTHONY WILLIAM VISSER

B.A., The University of Waterloo, 1995 M.A., The University of Waterloo, 1997

# A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

# DOCTOR OF PHILOSOPHY

in

# THE FACULTY OF GRADUATE STUDIES

Department of Psychology; Cognitive Systems

We accept this thesis as conforming to the required standard

# THE UNIVERSITY OF BRITISH COLUMBIA

July 2001

© Troy Anthony William Visser, 2001

In presenting this thesis in partial fulfilment of the requirements for an advanced degree at the University of British Columbia, I agree that the Library shall make it freely available for reference and study. I further agree that permission for extensive copying of this thesis for scholarly purposes may be granted by the head of my department or by his or her representatives. It is understood that copying or publication of this thesis for financial gain shall not be allowed without my written permission.

Department of Psychology

The University of British Columbia Vancouver, Canada

Date AUGUST 30, 2001

#### ABSTRACT

The present work is an examination of the mechanisms underlying temporal processing in vision. Recent studies have shown that when observers are asked to identify two objects presented in rapid succession, identification of the first object is quite accurate, while identification of the second object is poor when it follows the first at very brief inter-target intervals (i.e. 200-500 ms). This second-target deficit is known as the attentional blink. According to bottleneck models, the attentional blink occurs because processing of the first target prevents the second target from gaining access to high-level processing. A strong prediction of this account is that if processing time for the first target is increased, the magnitude of the attentional blink should also increase. This prediction is confirmed in five experiments. It is argued that these results strongly support bottleneck models as an account of the attentional blink in particular and of temporal processing more generally.

# TABLE OF CONTENTS

٤

Abstract ii	i
Table of Contents iii	ii
List of Figures i	v
Acknowledgementsv	i
General Introduction	
Experiment 1A	)
Method9Results and Discussion1	) .0
Experiment 1B	2
Method1Results and Discussion1	2
Experiment 2	21
Method2Results and Discussion2	22 24
Experiment 3	29
Method3Results and Discussion3	0 0
Experiment 4	2
Method3Results and Discussion3	13 13
Experiment 5	5
Method3Results and Discussion3	15 16
General Discussion	57

References	. 57
Figures	. 63
Appendix A	. 71

# LIST OF FIGURES

Figure 1.	Mean response times from Experiment 1A	63
Figure 2.	Representation of figures used as masks	64
Figure 3.	Mean accuracy from Experiment 1B	65
Figure 4.	Mean accuracy from Experiment 2	66
Figure 5.	Mean response times from Experiment 3	67
Figure 6.	Mean accuracy from from Experiment 3	68
Figure 7.	Mean response times from Experiment 4	69
Figure 8.	Mean T2-Mask ISIs from Experiment 5	70

#### ACKNOWLEDGEMENTS

This thesis would not have been possible without the love, friendship, and guidance of so many people in my life. Thank you to my parents, Terry and Nancy, who have supported and nurtured me through some quarter-century of "going to school". Thank you to Vince Di Lollo, and Jim Enns for showing me how to be a scientist, but more importantly to enjoy the mystery and the wonder of our craft. Finally, thank you to those special people who put up with my ravings, laugh at my jokes (sometimes), patiently advise me, and make sure that I'm still okay: Jeneva Ohan, David Prime, Catherine Spielmacher, Bryan Tham, and Samantha Zuvic. Many tasks require us to perceive rapid series of visual inputs. For example, to drive a car, we must attend to a stream of oncoming vehicles, road signs, pedestrians, and traffic lights in order to navigate and avoid collisions with other objects. Even a relatively simple task, such as finding a person in a crowd, requires rapid processing of a series of faces that impinge upon our retinas as we scan the environment. Of interest in the present work is an investigation of the mechanisms that underlie this type of temporal processing in vision.

Recent studies of temporal processing have employed a paradigm in which observers are presented with two target stimuli in rapid succession. Using a variety of different stimuli, such as letters (Raymond, Shapiro, & Arnell, 1992), color patches (Ross & Jolicoeur, 1999), and shapes (Joseph, Chun, & Nakayama, 1998), these studies have yielded a common finding. Observers are highly accurate when identifying an initial target (T1), but are poor at identifying a second target (T2) when it follows the first at a brief inter-target interval. This second-target deficit is often referred to as the <u>attentional blink</u> (AB; Broadbent & Broadbent, 1987; Raymond et al., 1992). The fact that the AB has been obtained using a variety of stimuli speaks to the generality of this deficit in visual processing. Moreover, that it occurs only when a preceding target must also be identified suggests that the AB is mediated by higher-level processing mechanisms (i.e. attention). However, although these points are generally agreed upon, the specific mechanisms underlying the AB remain debated.

One class of models that has been proposed to explain the AB is bottleneck models. According to these models, temporal processing mechanisms can be characterized in terms of two broadly sequential processing stages. In an initial parallel stage, attributes of incoming stimuli such as color, orientation and form are rapidly analyzed in order to identify potential targets. Then, in a second serial stage, representations of potential targets are coded into a form suitable for memory consolidation and response planning (e.g. Chun & Potter, 1995; Giesbrecht & Di Lollo, 1998; Jolicoeur, 1998). This architecture provides a parsimonious account of the basic AB phenomenon. However, a stronger test of this class of models can be made by determining how the amount of time it takes to process T1 affects the magnitude of the AB. As explained in greater detail below, bottleneck models predict that T1 processing time should be strongly related to the magnitude of the AB, such that an increase in T1 processing time should increase the magnitude of the AB.

To date, studies examining the relationship between T1 processing time and AB magnitude have done so <u>indirectly</u> by either correlating T1 accuracy with measures of AB magnitude (e.g. Seiffert & Di Lollo, 1997) or by examining the effect of T1-difficulty on the AB (Ward, Duncan, & Shapiro, 1997; McLaughlin, Shore, & Klein, 2001). The outcomes of these studies have been contradictory. Moreover, the use of indirect measures was based on the questionable assumption that the time taken to process T1 increases with T1 difficulty. For these reasons, in the present work, the relationship between T1 processing time and AB magnitude was examined <u>directly</u> by measuring both T1 processing time and AB magnitude. To anticipate the results, evidence in favour of bottleneck models was obtained consistently across five experiments.

#### The Attentional Blink

ŕ.

Two main paradigms have been used to study the AB. In one, known as rapid-serialvisual-presentation (RSVP), two targets are embedded in a stream of distractors presented at a central location (e.g. Raymond et al, 1992; Chun & Potter, 1995). The number of intervening distractors between T1 and T2 is varied systematically in order to examine the influence of inter-target interval (lag) on target identification. In a second method, referred to as the two-target paradigm (e.g. Duncan, Martens, & Ward, 1994; Ward et al., 1997), the targets are presented without a distractor stream. Instead, a single mask follows each target, and the inter-target lag is manipulated by inserting a variable temporal interval during which the display is blank between the targets. Additionally, rather than all items appearing in a central location, the two targets are presented in different eccentric locations.

Despite significant procedural differences between the RSVP and two-target paradigms, they have yielded homologous results. In both cases, identification of T1 is highly accurate, regardless of lag. In contrast, identification of T2 varies as a function of lag with poorest performance at shorter lags of 200-300 ms, and a gradual improvement to the level of T1 accuracy by a lag of 500-700 ms. This pattern of gradual improvement in T2 performance across lags is known as the attentional blink.

### Bottleneck Models of the Attentional Blink

The AB has generally been attributed to the requirement to attend to T1. Strong support for this view was obtained by Raymond et al. (1992) who demonstrated that the AB could be eliminated simply by instructing observers to ignore T1. This finding suggested that the mere presence of T1 was insufficient to impair T2 performance. Instead, it was necessary for observers explicitly to attend to T1, thereby engaging some sort of limited-capacity resource that was made unavailable for T2 when it followed T1 at a shorter lag.

Despite this broad consensus about the importance of attention in producing the AB, there is little agreement about the specific mechanisms involved. There are two main classes of theories about the AB – interference models (e.g. Shapiro, Raymond, & Arnell, 1994; Shapiro & Raymond, 1994) and bottleneck models (e.g. Chun & Potter, 1995; Jolicoeur, 1998). The implications of the present results for interference models are addressed in the General Discussion. First, however, the main tenets of bottleneck models are described below in some detail since the main focus of the present work is to evaluate the predictions of these models.

According to bottleneck models, temporal processing is said to involve two or more broadly sequential stages – an initial stage that is essentially parallel, followed by one or more subsequent stages that are capacity-limited and serial. A representative example is the two-stage model of Chun and Potter (1995). In the first stage, stimulus features such as colour, orientation, and form are rapidly analyzed in parallel across the visual field in order to detect potential targets. Although this analysis seems to be quite extensive, extending to the level of semantics (e.g. Luck, Vogel, & Shapiro, 1996; Maki, Frigen, & Paulson, 1997), the resulting representations of items are subject to rapid decay and interference from trailing stimuli. Therefore, representations in Stage 1 must be transferred to Stage 2 where they can be processed to a level appropriate for memory encoding, response planning, and execution.

Critically, this second stage is said to be serial and capacity-limited in that it can process only one item at a time. As a result, if T2 arrives while Stage 2 is busy with T1, it will be delayed in Stage 1 where it is vulnerable to decay or overwriting by subsequent stimuli (Chun & Potter, 1995; Giesbrecht & Di Lollo, 1998; Brehaut, Enns, & Di Lollo, 1999). The assumption that unattended objects are more vulnerable to masking is supported by a plethora of studies in the visual-masking literature (see Appendix A for a comprehensive review of this literature). Such studies indicate that the perception of unattended objects is more likely to be interfered with both by a mask presented in the same location as the target (e.g. Spencer & Shuntich, 1970; Ramachandran & Cobb, 1995), as well as by items presented in adjacent locations (Nazir, 1992).

According to bottleneck models, the AB is said to occur at shorter lags because T2 is more likely to arrive while Stage 2 is still busy processing T1. In turn, this increases the likelihood that T2 will be delayed in Stage 1 where it is subject to decay and masking. As lag increases, however, the AB declines in magnitude because it becomes increasingly likely that T1 processing will be completed before T2 is presented. This allows T2 to gain immediate access to Stage 2, thereby avoiding decay or masking.

#### Testing the Bottleneck Models

One unambiguous prediction of bottleneck models concerns the effect of T1 processing time on the magnitude of the AB. According to these models, increasing T1 processing time should produce a larger AB. This is because T1 will occupy Stage 2 for longer, thereby increasing the probability that T2 will be delayed in Stage 1 where it is subject to decay and masking.

To test this prediction, previous studies have used two indirect methods. In one set of studies, T1 accuracy was correlated with the magnitude of the AB. The logic behind this correlative analysis rested on the assumption that the time taken to process T1 increases directly with T1-difficulty. Given this assumption, it follows that bottleneck models must

predict that a decrease in T1 accuracy will be accompanied by an increase in the magnitude of the AB. Consistent with this prediction, Shapiro et al. (1994), Grandison, Ghirardelli, and Egeth (1997) and Seiffert and Di Lollo (1997) all reported negative correlations between T1 accuracy and AB magnitude.

A second set of studies tested bottleneck models by examining the effect of "T1difficulty" on the AB. In these studies, observers' ability to discriminate or identify T1 was varied systematically, and the effect of this variation on T2 performance was examined. The logic underlying this manipulation was similar to that used in the correlative analyses. Namely, it was assumed that by making T1 more difficult to identify, it would take longer to process. Given this assumption, bottleneck models must predict that an increase in T1difficulty would increase the magnitude of the AB because T2 would be delayed in Stage 1 for longer.

One representative study was that of Ward et al. (1997) who presented a pair of targets embedded in a stream of distractors, using a standard RSVP paradigm. Distractors were black letters. The first target was a white letter or a white rectangle. The second target was a black 'X' that was present on 50% of trials. Observers were required to judge the size of T1 and to determine whether T2 was present or absent. To vary T1-difficulty, the size-difference between a large and a small T1 was varied. In the "easy" condition, large and small targets were highly discriminable from one another; in the "hard" condition, large and small targets were much less discriminable from one another. Ward et al. (1997) found that T1 accuracy was significantly lower when the task was "hard" than when it was "easy". This

difference attested to the effectiveness of the size manipulation. Nonetheless, despite a substantial effect on T1 accuracy, level of difficulty had no effect on AB magnitude.

In a conceptually similar experiment, using a variation of the two-target paradigm, McLaughlin et al. (2001) presented a pair of central targets, each followed by a single pattern mask. To vary T1-difficulty, the duration of the target and the mask was varied reciprocally while holding constant the inter-stimulus interval (ISI) between T1 and its mask and the total duration of the display. For example, in the "easy" condition, T1 duration was 45 ms, the ISI was 15 ms, and the mask duration was 45 ms. In the "hard" condition, T1 duration was 15 ms, the ISI was 15 ms, and the mask duration was 75 ms. Attesting to the effectiveness of the difficulty manipulation, T1 accuracy was significantly lower in the "hard" than in the "easy" task. However, replicating Ward et al. (1997), no influence of T1-difficulty was found on the magnitude of the AB.

Taken together, the results of Ward et al. (1997) and McLaughlin et al. (2001) seem to disconfirm the predictions of bottleneck models. However, before arriving at this conclusion, it is important to consider two potential problems. First, although the purpose of both the correlative and experimental studies noted above was to examine the relationship between T1 processing time and the AB, T1 processing time was never measured explicitly. Rather, it was simply assumed that a decrease in T1 accuracy indexed an increase in T1 processing time. Such a relationship, although plausible, is by no means assured. For example, in the case of "speed-accuracy" tradeoff, a <u>decrease</u> in target accuracy is accompanied by a corresponding decrease in target processing time. Given this consideration, it must be concluded that an inverse relationship between accuracy and processing time cannot merely be assumed: it must be demonstrated.

A second potential problem arises from the use of a mask after T1. In the experiments of Ward et al. (1997) and McLaughlin et al. (2001), a pattern mask was always presented 100 ms after the onset of T1, regardless of the level of T1-difficulty. A number of previous studies (e.g. Turvey, 1973; Breitmeyer & Ganz, 1976; Michaels & Turvey, 1979; Breitmeyer, 1984) have suggested that under these circumstances, a mask can terminate target processing. That is, the mask can pre-empt processing resources originally allocated to the target (see Appendix A). In the context of the T1-difficulty experiments, this raises the possibility that although T1 processing time may have been longer when T1-difficulty was greater, this increase was nullified by the onset of the trailing mask. In turn, if the mask nullified differences in T1 processing time, then according to bottleneck models, there should be no difference in AB magnitude.

To illustrate this possibility, consider the following example. Assume that in the "easy" task, 150 ms of processing time was required for accurate performance while in the "hard" task, 300 ms of processing time was necessary. According to bottleneck models, the additional 150 ms of processing time required in the "hard" task should produce a much larger AB. However, if a mask presented 100 ms after T1 interrupted its processing, this would limited the total duration of T1 processing to 100 ms in both the "hard" and "easy" tasks. As a result, according to bottleneck models, the magnitude of the AB would be equated across levels of T1-difficulty.

In summary, previous experiments that have examined the relationship between T1 processing time and AB magnitude share two potential problems. First, changes in T1 processing time were never assessed directly. Second, the role of the T1 mask in modulating the relationship between T1 processing time and AB magnitude was never determined. To address these concerns, it was first necessary to find a task that demonstrably affected T1 processing time. This was done in Experiment 1A by measuring response times to perform a size-discrimination task similar to that of Ward et al. (1997). Then having shown direct evidence for changes in T1 processing time using this task, it was used in Experiment 1B as the first of two tasks in a modified two-target AB paradigm. The purpose of this additional experiment was to examine the role of masking in modulating the relationship between T1 processing time and AB magnitude. Two conditions were employed. In one, T1 was masked and in the other, the mask was omitted.

#### **EXPERIMENT 1A**

The purpose of Experiment 1A was to estimate T1 processing time in a sizediscrimination task similar to that of Ward et al. (1997). To do this, observers were presented with a single large or small outline rectangle on each trial and were asked to make a speeded judgment about its size. In one case, the large and small rectangles were very similar in size. This was expected to result in relatively long processing times and relatively slow response times. In the other case, the large and small rectangles were very different in size. This was expected to produce relatively short processing times and relatively fast response times.

# Method

<u>Participants</u>. Sixteen undergraduate students (8 female) at the University of British Columbia participated for course credit. All participants reported normal or corrected-tonormal vision.

<u>Apparatus and Stimuli</u>. All stimuli had a luminance of 10 cd/m<sup>2</sup>, as measured by a Minolta LS-100 luminance meter, and were displayed on a Tektronix 608 oscilloscope, equipped with fast P15 phosphor. The viewing distance, set by a headrest, was 57 cm. The background and surrounding visual field were dark, except for dim illumination of the keyboard. The target was an outline rectangle shape that could be one of three sizes:  $0.50^{\circ}$  (width) x  $1.0^{\circ}$  (height),  $0.31^{\circ}$  x  $0.63^{\circ}$ , or  $0.25^{\circ}$  x  $0.50^{\circ}$ .

<u>Procedure</u>. There were two blocks of 80 trials. In the "easy" block, observers were presented with either the largest or the smallest rectangle on each trial. In the "hard" block, observers were presented with either the largest, or the mid-sized rectangle on each trial. The order in which observers received each block was counterbalanced, such that half of the observers did the "easy" block first. Each block of trials began with a display in which the "large" and "small" rectangles were shown side-by-side. Observers were instructed to memorize the size of each rectangle in order to be able to discriminate between them accurately and then to begin the experimental trials by pressing the space bar.

Each trial began with the presentation of a fixation point in the centre of the screen. Observers were instructed to focus on the fixation dot and then to press the space bar to start the sequence of stimuli. Following a blank screen that lasted for a random interval between 500 and 800 ms, the target rectangle was presented in the center of the screen for 30 ms. Observers were instructed to indicate as quickly as possible whether the T1 rectangle was either "large" or "small", by pressing one of two appropriately-marked keys on the keyboard. After making this response, the next trial began with the presentation of a central fixation point.

#### Results

Trials on which errors were made were discarded from the response-time analysis. This amounted to a total of 4.3% of trials in the "easy" block and 8.9% in the "hard" block. Response times from all other trials were screened using the outlier procedure described by Van Selst and Jolicoeur (1994), which employs a floating criterion based on sample size to determine outliers. This resulted in the removal of a further 2.2% of RTs from the "easy" block, and 1.7% from the "hard" block. The remaining data were used to calculate mean RTs for each block. These means are illustrated in Figure 1. A paired-samples t-test indicated that

Insert Figure 1 about here

\_\_\_\_\_

mean RTs in the "hard" block were significantly longer than those in the "easy" block,  $\underline{t}(15) = 6.65$ ,  $\underline{p} < .001$ , d = 1.66. Importantly, because this increase in response times was accompanied by an increase in error rate, interpretation of the results is not compromised by a speed-accuracy tradeoff.

### Discussion

The present results indicate that response times were significantly longer when the size-discrimination was more difficult. This suggests that target processing time was longer in the more difficult task, thus supporting the assumption made in earlier studies that T1 processing time increases directly with T1-difficulty (e.g. Ward et al., 1997). Having obtained direct evidence for changes in T1 processing time in the size-discrimination task, what remains to be explained is why Ward et al. (1997) failed to obtain commensurate changes in the magnitude of the AB. To address this issue, in the next experiment, we investigate the role of the T1 mask in order to determine the relationship between T1 processing time and AB magnitude.

#### **EXPERIMENT 1B**

Experiment 1B had two goals. The first was to determine whether T1 processing time would modulate the magnitude of the AB. The second was to determine whether this relationship would be influenced by the presence of a mask after T1. On each trial, observers were presented with two sequential targets, using a variation of the two-target paradigm. The first target was a rectangle identical to those used in Experiment 1A. The second was a letter to be identified. This letter was always followed by a pattern mask.

To examine the influence of T1 processing time, observers received two blocks of trials in which the difficulty of the size-discrimination task used in Experiment 1A was varied. Given the findings from Experiment 1A, it was expected that T1 processing time would be relatively short when the size-discrimination was simple, and relatively long when it was difficult. To examine the role of the mask, T1 was always followed by a mask for one

group of observers, and never followed by a mask for another group. This yielded a  $2 \ge 2$  design with T1 processing time as a within-subjects factor, and T1 masking as a between-subjects factor.

# Method

Participants. Thirty-six undergraduate students (21 female) at the University of British Columbia participated for course credit. All participants reported normal or corrected-to-normal vision. Twenty-four students participated in the condition with a mask after T1, and twelve participated in the condition without a mask after T1.

Apparatus and Stimuli. The apparatus, stimuli and viewing conditions were identical to those employed in Experiment 1A. The first target was an outline rectangle shape that could be one of three sizes:  $0.50^{\circ}$  (width) x  $1.0^{\circ}$  (height),  $0.31^{\circ}$  x  $0.63^{\circ}$ , or  $0.25^{\circ}$  x  $0.50^{\circ}$ . The second target was any letter of the English alphabet except for I, O, Q, and Z. These four letters were omitted on the grounds that they were confusable with digits. The masks used for targets consisted of geometric shapes, constructed from letter segments. Although the masks shared visual features with letters, they were not confusable with any English letter. An example of these masks can be seen in Figure 2. Both the second target and the masks subtended an area of approximately  $1^{\circ}$  square of visual angle.

Insert Figure 2 about here

\_\_\_\_\_

\_\_\_\_\_

<u>Procedure</u>. There were two conditions – one in which T1 was masked and a second . in which T1 was never masked. Within each condition, there were two blocks of trials. In one block, now referred to as "short" because T1 processing time was expected to be relatively short, observers were presented with either the largest or the smallest rectangle as T1 on each trial. In the other block, now referred to as "long" because T1 processing time was expected to be relatively long, observers were presented with either the largest, or the mid-sized rectangle as T1 on each trial. In both conditions, the second target was a single letter, chosen randomly with replacement. This second target was always masked. Each block of trials began with a display in which the "large" and "small" T1 rectangles were displayed side-by-side. Observers were told to memorize the size of each rectangle in order to be able to discriminate between them accurately and then to begin the experimental trials by pressing the space bar.

Each trial began with the presentation of a fixation point in the centre of the screen. Observers were instructed to focus on the fixation dot and then to press the space bar to start the sequence of stimuli. Following a random interval between 500-800 ms during which the screen was blank, T1 was presented in the centre of the screen for 30 ms. In the condition in which T1 was masked, T1 was followed by an ISI of 70 ms during which the screen was blank, and then by the presentation of a mask for 30 ms, yielding a T1-mask stimulus-onsetasynchrony (SOA) of 100 ms. In the condition in which T1 was not masked, no intervening items separated T1 and T2.

The second target followed T1 at one of four lags corresponding to T1-T2 SOAs of 200, 300, 500, or 700 ms. It remained on the display for 30 ms and was followed by a 70 ms ISI during which the screen was blank, and then by a pattern mask that remained on the screen for 30 ms. This yielded a T2-mask SOA of 100 ms. Observers were instructed to

make two responses after the offset of the mask. The first was to indicate whether the T1 rectangle was either "large" or "small", by pressing one of two appropriately-marked keys on the keyboard. The second was to identify the T2 letter by typing it into the keyboard. After making these two responses, the next trial began with the presentation of a central fixation point.

Observers participated in only one of the two conditions. However, within their respective condition, all observers received both the "short" and "long" blocks. The order of these blocks was counterbalanced, such that half the observers received the "short" block of trials first. Each block began with 20 practice trials, followed by 80 experimental trials and consisted of equal numbers of trials with "large" and "small" rectangles. Experimental trials were divided evenly across T1-T2 lags yielding 20 trials at each of the four lags.

#### Results

<u>T1 Results.</u> With a mask after T1, mean percentages of correct identifications of T1 were 96.3 and 77.0 for the "short" and "long" blocks, respectively. Without a mask after T1, mean percentages were 99.6 and 94.5, respectively. These results were analyzed in a 2 (Processing Time: Slow vs. Fast) x 4 (T1-T2 Lag: 200, 300, 500 and 700 ms) x 2 (Condition: Mask vs. No Mask) mixed analysis of variance with Condition as a between-subjects factor. This analysis revealed a significant effect of Processing-time,  $\underline{F}(1, 34) = 49.98$ , p < .001, MS<sub>e</sub> = 189.39,  $\eta^2 = .60$ , as well as a significant interaction between Processing-time and Condition,  $\underline{F}(1, 34) = .16.83$ , p < .001, MS<sub>e</sub> = 189.39,  $\eta^2 = .33$ , indicating that the difference in T1 accuracy between the "short" and "long" blocks was significantly larger when T1 was masked than when the mask was omitted. <u>T2 Results.</u> Estimates of T2-identification were based exclusively on those trials in which T1 had been identified correctly. Mean percentages of correct T2 identification as a function of T1-processing time and T1-T2 lag, are illustrated in Figure 3 for both the condition in which the T1 mask was present (dotted lines), and the condition in which it was

\_\_\_\_\_

### Insert Figure 3 about here

\_\_\_\_\_

omitted (solid lines). Three aspects of these results are notable. First, in both conditions, T2 performance improved rapidly as T1-T2 lag is increased, indicating a robust AB. Second, collapsing across T1 processing time, the magnitude of the AB was smaller when the mask after T1 was omitted. Third, and most importantly, while the magnitude of the AB did not appear to vary as a function of T1 processing time when T1 was masked, such a relationship was apparent when the mask after T1 was omitted.

To confirm these impressions, the data were analyzed in a 2 (Processing Time: Slow vs. Fast) x 4 (T1-T2 Lag: 200, 300, 500 and 700 ms) x 2 (Condition: T1 Mask vs. No T1 Mask) mixed analysis of variance with Condition as a between-subjects factor. This analysis yielded a significant effect of Lag,  $\underline{F}(3, 102) = 69.62$ ,  $\underline{p} < .001$ , MS<sub>e</sub> = 106.06,  $\eta^2 = .67$ , indicating that a robust AB was obtained whether T1 was masked or not. Additionally, the results indicated a significant effect of Condition,  $\underline{F}(1, 34) = 18.65$ ,  $\underline{p} < .001$ , MS<sub>e</sub> = 445.38,  $\eta^2 = .35$ , and a Condition x Lag interaction,  $\underline{F}(3, 102) = 9.30$ ,  $\underline{p} < .001$ , MS<sub>e</sub> = 106.06,  $\eta^2 = .22$ , indicating that the overall magnitude of the AB was smaller when the mask was omitted.

### Bottleneck Models of Temporal Processing 17

Finally, separate 2 (Processing Time) x 4 (T1-T2 Lag) analyses for each condition revealed a significant interaction when there was no T1 mask,  $\underline{F}(3, 33) = 5.04$ ,  $\underline{p} = .006$ , MS<sub>e</sub> = 32.26,  $\eta^2 = .31$ , but no interaction when there was a T1 mask,  $\underline{F}(3, 69) = 1.67$ ,  $\underline{p} = .18$ , MS<sub>e</sub> = 99.80,  $\eta^2 = .07$ . This pattern of results is consistent with the hypothesis that T1 processing time modulates AB magnitude when there is no mask after T1.

Order and Practice Effects. In order to clarify the nature of the effects found in Experiment 1B, an additional analysis was performed to determine whether the order in which observers received the "short" or "long" blocks influenced T1 or T2 performance. For the condition in which T1 was masked, both T1 and T2 scores were submitted to a 2 (Processing Time) x 4 (T1-T2 Lag) x 2 (Order: "Slow" Block First vs. "Fast" Block First) mixed-design analysis of variance with Order as a between-subjects factor.

For T1 scores, this analysis revealed a significant main effect of Order,  $\underline{F}(1, 22) = 5.52$ ,  $\underline{p} = .028$ ,  $MS_e = 204.15$ ,  $\eta^2 = .20$ , and an interaction between Processing Time and Order,  $\underline{F}(1, 22) = 6.64$ ,  $\underline{p} = .017$ ,  $MS_e = 216.23$ ,  $\eta^2 = .23$ . An inspection of the data suggests that while T1 accuracy was similar in the "short" block regardless of which block came first, accuracy was much higher in the "long" block when it was presented second. This suggests that performance on the relatively hard size-discrimination task benefited from a prior block of trials with the relatively easy size-judgment. For T2 scores, however, no effects involving Order were significant (all p's > .34,  $\eta^{2'}s < .05$ ). This suggests that although the order in which observers saw each block of trials influenced T1 scores, there was no evidence of order effects on the magnitude of the AB.

A similar analysis of T1 and T2 scores was performed for the condition in which T1 was not masked. For T1 scores, no effects involving Order were significant (all p's > .47,  $\eta^2$ 's < .08). Similarly, there were no significant effects involving Order for T2 (all p's > .42,  $\eta^2$ 's < .09). Collectively, these results suggest that the order in which observers received the "short" and "long" blocks did little to influence accuracy. Most importantly, there was no evidence that order influenced the magnitude of the AB.

A final analysis was performed on T1 and T2 accuracy data from the condition in which the mask after T1 was omitted. The purpose of this analysis was to determine whether practice within a block of trials would influence accuracy by comparing scores from the first half of trials with those from the second half. For the "short" block, both T1 and T2 scores were analyzed in a 2 (Half: first vs. second half of trials) x 4 (T1-T2 Lag) x 2 (Order: "Short" Block First vs. "Long" Block First) mixed-design analysis of variance with Order as a between-subjects factor. The analysis revealed no significant effects of Half for either T1 accuracy (all p's > .24,  $\eta^2$ 's < .17) or T2 accuracy (all p's > .66,  $\eta^2$ 's < .13). A similar analysis for the "long" block also showed no significant effects of Half for either T1 accuracy (all p's > .08,  $\eta^2$ 's < .27) or T2 accuracy (all p's > .24,  $\eta^2$ 's < .16). Collectively, these results suggest that there was no influence of practice within a block of trials on either T1 performance or AB magnitude.

# Discussion

In Experiment 1B, a strong relationship between T1 processing time and AB magnitude was obtained when T1 was not masked, but not when a mask was presented after T1. Two conclusions are invited by these findings. First, the fact that T1 processing time

modulated AB magnitude when T1 was not masked strongly supports the predictions of bottleneck models. According to these models, the magnitude of the AB should vary directly as a function of the time required to process T1. This is because the duration of T1 processing determines how long T2 remains delayed in Stage 1 and thus vulnerable to masking. These results are also consistent with other studies that have found that unattended targets are more vulnerable to masking by a temporally-trailing item (see Appendix A).

Second, the failure to find a relationship between T1 processing duration and AB magnitude when T1 was masked provides some insight into previous results obtained by Ward et al. (1997) and McLaughlin et al. (2001) who found no relationship between T1-difficulty and AB magnitude. The present results suggest that this outcome can be attributed to the presence of a mask after T1, which interrupted T1 processing. Because T1-mask SOA was the same in both the "easy" and "hard" conditions in these experiments, this interruption effectively equated T1 processing time across levels of T1-difficulty. Given equivalent processing times, bottleneck models predict that the magnitude of the AB would also be the same.

Additional support for this account comes from an analysis of T1 error rates in the present experiment as well as those of Ward et al. (1997) and McLaughlin et al. (2001). In every case, with a mask after T1, error rates were almost 20% higher when the T1 task was difficult (or equivalently, when T1 processing took a long time). This can be contrasted to the case in which the mask after T1 was omitted. Under these conditions, the effect of difficulty on T1 accuracy was much smaller.

A simple account of these findings can be made in terms of interruption masking. Consider first the large differences in T1 accuracy found when T1 was masked. This difference can be attributed to the fact that interruption of processing becomes increasingly deleterious to accuracy as the total processing time required to identify a target increases. For example, if T1 requires 150 ms to identify, but processing is interrupted by the mask after 100 ms, accuracy would be impaired relatively little. This is because much of the necessary processing would have been completed before the mask was presented. In contrast, if T1 requires 300 ms to identify, but the mask interrupts processing after 100 ms, a relatively large impairment would occur. This is because very little of the required processing would be completed before the mask was presented. A similar explanation can be advanced to explain why T1 accuracy improves so markedly when the mask is omitted. Under these conditions, T1 processing can continue uninterrupted. As a result, even when the processing time required to identify T1 is relatively long, there is a high probability that processing will be completed successfully.

One unexpected aspect of the results was that regardless of T1 processing time T2 accuracy was much better when the mask after T1 was omitted. An account of this difference can be made in terms of integration masking. Turvey (1973) noted that when targets and masks are separated by relatively short SOAs, target degradation occurs via a process of contour integration whereby the contours of the target and the mask combine in a process akin to the addition of noise (the mask) to a signal (the target). Importantly, a single mask can produce both integration and interruption masking (Michaels & Turvey, 1979). Thus, it is possible that the mask after T1 not only increased the total processing time

required to identify T1 relative to the case in which the mask after T1 was omitted, but also interrupted T1 processing. If this were the case, then according to bottleneck models, overall T2 performance would be more impaired with a mask after T1 than when the mask was omitted.

Although the present results suggest that T1 processing time is correlated with AB magnitude, it is desirable to test this relationship over a broader range of T1 processing times in order to verify these findings. One paradigm appropriate for such an investigation is visual search. In a typical visual search experiment, a target must be picked out from amongst a varying number of simultaneously-presented distractor items. The time it takes to find the target typically varies as a function of the number of distractors. With few distractors, it takes very little time to find the target. With many distractors, it takes quite a long time. Thus, by varying systematically the number of distractors, target-processing time can be varied across an extensive range.

In Experiment 2, a visual search task was used to vary T1 processing time. Observers were instructed to find a designated target letter presented amongst a variable number of distractor letters. The number of distractors was varied systematically in order to manipulate the time required to find T1. In accordance with the predictions of bottleneck models, a strong relationship was expected between set-size (i.e. the total number of items in the search display) and AB magnitude.

#### **EXPERIMENT 2**

The purpose of Experiment 2 was to replicate and extend the findings from Experiment 1 demonstrating that T1 processing time modulates AB magnitude. On each trial, observers were presented with two targets. The first target was the letter 'C' or 'G', presented either alone, or with one, or four distractor letters. The second target was a single letter to be identified. Observers were asked to decide whether T1 was a 'C' or a 'G' and to identify T2.

There were two conditions. In one, the T1 search display was masked by a trailing display of digits. In the other condition, the search display was never masked. Based on the results of Experiment 1B, it was expected that AB magnitude would <u>not</u> vary as a function of set-size when the search display was masked. This would be consistent with the hypothesis that the mask interrupts T1 processing. In contrast, with no mask after T1, AB magnitude was expected to increase as a function of set-size, reflecting an increase in the time required to find T1.

#### <u>Method</u>

<u>Participants</u>. Thirty-six undergraduate students (20 female) at the University of British Columbia participated for course credit. All participants reported normal or corrected-to-normal vision. None had participated in the previous experiments. Twenty-four students participated in the condition in which T1 was masked, while the remaining twelve participated in the condition in which T1 was not masked.

<u>Apparatus and Stimuli</u>. The apparatus and viewing conditions were identical to those in Experiment 1B. In the first display, T1 was always either the letter 'C' or 'G', while distractors could be any letter in the English alphabet except for I, O, Q, Z, C, or G. These letters were omitted on the grounds that they were either confusable with digits (i.e. I, O, Q, Z) or were identical to the first target (i.e. C, G). The second target was any letter of the

0

English alphabet except for I, O, Q, Z, C, or G. The same criterion for choosing T2 was used as for choosing distractors in the first display. Masks consisted of digits from one to nine. The identical digit was used to mask both T1 and T2. Targets, distractors, and masks all subtended an area of approximately 1° square of visual angle.

<u>Procedure</u>. There were two conditions in Experiment 2. In one condition, the T1 search display was followed by a mask display; in the other, the T1 display was never masked. Each condition was divided into three blocks in which the number of distractors presented along with the target numbered either zero (set-size 1), one (set-size 2), or four (set-size 5). For the purposes of creating the search display, the screen was divided into a notional 5 x 5 matrix of locations. The target and any accompanying distractors could appear at any location in this matrix with two constraints. First, no item could appear in the centre location. Second, items could not be directly adjacent to one another.

Each trial began with the presentation of a fixation point in the centre of the screen. Observers were instructed to press the space bar to start the sequence of stimuli. Following a random interval from 500 to 800 ms during which the screen was blank, the T1 search display was presented for 30 ms. In the condition in which T1 was masked, this search display was followed by an ISI of 70 ms during which the screen was blank, and then by a 30 ms presentation of a mask display. This display consisted of a digit presented at each location where a letter had been in the search display. The same digit was presented in all locations. In the condition in which T1 was not masked, there were no intervening items presented between T1 and T2. The second target, which was a letter, was presented in the centre of the screen and followed T1 at one of four lags corresponding to T1-T2 SOAs of 200, 300, 500, or 700 ms. The target remained on the screen for 30 ms and was followed, after a 70 ms ISI during which the screen was blank, by a pattern mask that remained on the screen for 30 ms. Observers were instructed to make two responses. The first response was always to indicate whether T1 was either a "C" or a "G", by pressing one of two appropriately-marked keys on the keyboard. Then, observers identified the T2 letter by typing it into the keyboard. After making these two responses, the next trial began with the presentation of the central fixation point.

Within in each condition, the order of the blocks was completely counterbalanced. Each block began with 10 practice trials, followed by 120 experimental trials. These experimental trials consisted of equal numbers of trials with "C" or "G" as T1. Trials were divided evenly across T1-T2 lags yielding 30 trials at each of the four lags.

#### Results

<u>T1 Masked.</u> Mean percentages of correct identifications of T1 were 95.1, 81.9 and 65.7 for the set-size 1, set-size 2, and set-size 5 blocks, respectively. The results were analyzed in a 3 (Set-size: 1, 2, 5) x 4 (T1-T2 Lag: 200, 300, 500 and 700 ms) repeated-measures analysis of variance. The analysis revealed only a significant effect of Set-size,  $\underline{F}(2, 46) = 94.49$ , p < .001, MS<sub>e</sub> = 220.61,  $\eta^2 = .80$ , confirming that overall T1 accuracy declined significantly as the number of distractors presented with T1 increased.

Estimates of T2-identification were based exclusively on trials in which T1 had been identified correctly. Mean percentages of correct T2 identification as a function of set-size

and T1-T2 lag are illustrated in Panel A of Figure 4. An inspection of this figure suggests that

### Insert Figure 4 about here

\_\_\_\_\_

#### \_\_\_\_\_

an AB occurred at all set-sizes. Moreover, it appears that while the AB was clearly smallest at set-size 1, there was no difference in AB magnitude at set-sizes 2 and 5. To confirm these impressions, the data were first analyzed in a 3 (Set-size) x 4 (T1-T2 Lag) repeated-measures analysis of variance. Confirming the presence of an AB, the analysis indicated a significant effect of Lag,  $\underline{F}(3, 69) = 97.44$ , p < .001,  $MS_e = 181.75$ ,  $\eta^2 = .81$ . In addition, there was a significant main effect of Set-size,  $\underline{F}(2, 46) = 29.30$ , p < .001,  $MS_e = 417.86$ , ,  $\eta^2 = .56$ , and a significant interaction between Lag and Set-size,  $\underline{F}(6, 138) = 28.11$ , p < .001,  $MS_e = 71.54$ , ,  $\eta^2 = .55$ , indicating that AB magnitude varied directly with set-size. In order to determine whether AB magnitude differed between set-sizes 2 and 5, the analysis was repeated using only these two set-sizes. Once again, there was a significant effect of Lag,  $\underline{F}(3, 69) = 100.76$ , p < .001,  $MS_e = 208.77$ ,  $\eta^2 = .81$ , indicating the presence of an AB. However, consistent with the impression that the magnitude of the AB did not differ between set-sizes 2 and 5, there was no significant effect of Set-size (p = .81,  $\eta^2 = .01$ ) or a Set-size x Lag interaction (p = .74,  $\eta^2 = .02$ ).

<u>No T1 Mask.</u> Mean percentages of correct identifications of T1 were 98.2, 95.8 and 95.5 for set-sizes 1, 2, and 5 respectively. The results were analyzed in a 3 (Set-size) x 4 (T1-T2 Lag) repeated-measures analysis of variance. The analysis revealed a significant

effect of Set-size,  $\underline{F}(2, 22) = 5.42$ ,  $\underline{p} = .012$ ,  $MS_e = 19.22$ ,  $\eta^2 = .33$ , confirming that overall T1 accuracy differed significantly as a function of the number of distractors. In addition, a significant effect of Lag was also obtained,  $\underline{F}(3, 33) = 3.68$ ,  $\underline{p} = .022$ ,  $MS_e = 20.75$ ,  $\eta^2 = .25$ . An examination of T1 accuracy scores suggests that this effect arose from an overall improvement in T1 accuracy as T1-T2 lag increased.

Although significant, the differences in T1 accuracy as a function of set-size appeared to be smaller than those found with a mask after T1. To verify this impression, T1 accuracy scores with and without a T1 mask were analyzed in a 3 (Set-size) x 4 (T1-T2 Lag) x 2 (Condition: T1-Masked vs. No T1-Mask) mixed analysis of variance with Condition as a between-subjects factor. The analysis indicated a significant interaction between Set-size and Condition,  $\underline{F}(2, 68) = 37.15$ , p < .001, MS<sub>e</sub> = 155.46,  $\eta^2 = .52$ , confirming that differences in T1 accuracy between set-sizes were significantly smaller with no mask presented after T1.

Estimates of T2-identification were based exclusively on those trials in which T1 had been identified correctly. Mean percentages of correct T2 identification as a function of T1 processing time and T1-T2 lag are illustrated in Panel B of Figure 4. An inspection of this figure suggests that an AB occurred at all set-sizes. Moreover, the magnitude of the AB varied directly as a function of set-size, with progressively larger ABs occurring as set-size increased. To confirm these impressions, the data were analyzed in a 3 (Set-size) x 4 (T1-T2 Lag) repeated-measures analysis of variance. Confirming the presence of an AB, the analysis indicated a significant effect of Lag,  $\underline{F}(3, 33) = 123.08$ , p < .001,  $MS_e = 68.58$ ,  $\eta^2 =$ .92. In addition, confirming the impression that AB magnitude varied with set-size, there was a significant main effect of Set-size,  $\underline{F}(2, 22) = 56.47$ ,  $\underline{p} < .001$ ,  $MS_e = 279.78$ ,  $\eta^2 = .84$ , and a significant interaction between Set-size and Lag,  $\underline{F}(6, 66) = 18.34$ ,  $\underline{p} < .001$ ,  $MS_e = 75.97$ ,  $\eta^2 = .63$ .

In order to determine whether AB magnitude differed between set-sizes 2 and 5, the analysis was repeated using only these two set-sizes. Once again, there was a significant effect of Lag,  $\underline{F}(3, 33) = 99.86$ ,  $\underline{p} < .001$ ,  $MS_e = 107.09$ ,  $\eta^2 = .90$ , indicating the presence of an AB. In addition, unlike the condition in which T1 was masked, there was a significant effect of Set-size,  $\underline{F}(1, 11) = 15.21$ ,  $\underline{p} < .002$ ,  $MS_e = 293.20$ ,  $\eta^2 = .58$  and a Set-size x Lag interaction,  $\underline{F}(3, 33) = 3.58$ ,  $\underline{p} = .024$ ,  $MS_e = 93.84$ ,  $\eta^2 = .25$ . Together, these results indicate that the magnitude of the AB was significantly larger at set-size 5 than at set-size 2.

# Discussion

The goal of Experiment 2 was to obtain further evidence that T1 processing time modulates AB magnitude. Consistent with this expectation, without a mask after T1, AB magnitude increased significantly as set-size was increased. This supports the prediction of bottleneck models that T1 processing time determines the magnitude of the AB deficit and replicates other studies that have shown unattended targets are more vulnerable to masking (see Appendix A). In addition, when T1 was masked, there was no difference in AB magnitude between set-sizes 2 and 5. This is consistent with findings in Experiment 1B showing that when T1 was masked, there was no difference in AB magnitude as a function of the difficulty of the size-discrimination task. Also consistent with the results of Experiment 1B, T1 accuracy decreased substantially with an increase in set-size when T1 was masked. However, the magnitude of this reduction was much decreased when the mask after T1 was omitted.

One puzzling finding is that a reliably smaller AB was obtained at set-size 1, even when the first target was masked. An explanation for this finding may rest on the fact that the time required to process T1 at set-sizes 2 and 5 was greater than the T1-mask SOA, while at set-size 1 it was less than the T1-mask SOA. If that were the case then the mask would prematurely terminate processing of T1 at set-sizes 2 and 5, thereby equating processing duration. However, at set-size 1, the mask would not influence processing duration at all. Under these circumstances, according to bottleneck models, the magnitude of the AB should be similar at set-sizes 2 and 5, but significantly smaller at set-size 1. This, of course, is exactly the pattern of results obtained in Experiment 2.

To illustrate this argument, consider the following example. Assume that at set-size 2, finding the target requires 150 ms of processing time, while at set-size 5 it requires 350 ms of processing time. Further, assume that the T1-mask SOA is 100 ms. Under these conditions, because the time required to find T1 in both conditions is longer than the T1-mask SOA, the mask would equate actual T1 processing times at 100 ms. Now, assume that a set-size 1 condition is added in which only 60 ms is required to find the target. Because processing time in this task is less than the T1-mask SOA, the mask would not influence processing duration. Therefore, actual T1 processing time would be 60 ms at set-size 1, 100 ms at set-size 2, and 100 ms at set-size 5. Given these processing times, bottleneck models must predict that the magnitude of the AB would be the same at set-sizes 2 and 5 because T1 processing time is the same, but smaller at set-size 1 because T1 processing time is shorter.

In the present experiment, it was assumed that the visual search task modulated T1 processing time on the basis of results obtained in prior visual search experiments (e.g. Treisman & Gelade, 1980; Logan, 1994). However, despite this evidence, it is important to demonstrate that a set-size manipulation has a similar effect in the context of an AB paradigm. To this end, in Experiment 3, the present experiment was replicated with the addition of an explicit measure of T1 processing time.

#### EXPERIMENT 3

Experiment 3 had two objectives. The first was to provide direct evidence, within the context of an AB paradigm, that varying set-size in the T1 visual search task modulated T1 processing time. The second was to replicate results from previous experiments indicating that T1 processing time and AB magnitude are positively correlated in the absence of a mask after T1. As in Experiment 2, T1 was a visual-search task and T2 was a letter-identification task. The key difference was that in order to obtain an estimate of T1 processing time, observers were asked to determine whether the search display contained a 'C' or 'G' as quickly as possible.

# Method

<u>Participants</u>. Sixteen undergraduate students (12 female) at the University of British Columbia participated for course credit. All participants reported normal or corrected-tonormal vision. None had participated in the previous experiments.

<u>Apparatus and Stimuli</u>. The apparatus and stimuli were identical to those used in Experiment 2.
<u>Procedure</u>. Stimulus presentation was identical to that used in the condition without a mask after T1 in Experiment 2, with one exception. On each trial, observers made a speeded response to T1, as soon as they could determine whether the search display contained a "C" or a "G". As a result, responses to T1 were often made while the display sequence on a trial was ongoing. This differed from Experiment 2 in which all responses were made after the display sequence on each trial had ended.

### Results

Mean percentages of correct identifications of T1 were 80.2, 77.0 and 71.1 for setsizes 1, 2, and 5 respectively. The results were analyzed in a 3 (Set-size: 1, 2, 5) x 4 (T1-T2 Lag: 200, 300, 500 and 700 ms) repeated-measures analysis of variance. The analysis revealed only a significant effect of Set-size,  $\underline{F}(2, 30) = 36.41$ , p < .001,  $MS_e = 37.62$ ,  $\eta^2 =$ .71, confirming that overall T1 accuracy decreased significantly as set-size increased.

Trials on which errors were made were discarded from the response-time analysis. Response times from all other trials were screened using the outlier procedure described by Van Selst and Jolicoeur (1994). This resulted in the removal of a further 3.4% of RTs from the set-size 1 block, 4.2% of RTs from the set-size 2 block, and 4.5% from the set-size 5 block. The remaining RT data were used to calculate mean RTs for each block. These means are illustrated in Figure 5. An inspection of this figure suggests that RTs increased

\_\_\_\_\_

Insert Figure 5 about here

\_\_\_\_\_

systematically with an increase in set-size. To confirm this impression, the results were analyzed in a 3 (Set-size) x 4 (T1-T2 Lag) repeated-measures analysis of variance. The analysis revealed only a significant effect of Set-size,  $\underline{F}(2, 30) = 13.33$ , p < .001, MS<sub>e</sub> = 21902.03,  $\eta^2 = .47$ , verifying that mean T1 RTs increased significantly as set-size increased. Importantly, because this increase in RTs was accompanied by an increase in error rate, interpretation of the results is not compromised by a speed-accuracy tradeoff.

Estimates of T2-identification were based exclusively on those trials in which T1 had been identified correctly. Mean percentages of correct T2 identification as a function of setsize and T1-T2 lag are illustrated in Figure 6. An inspection of this figure suggests that

## Insert Figure 6 about here

\_\_\_\_\_

an AB occurred at all T1 set-sizes. Moreover, the magnitude of the AB varied directly as a function of set-size, with a progressively larger AB as set-size increased. To confirm these impressions, the data were analyzed in a 3 (Set-size) x 4 (T1-T2 Lag) repeated-measures analysis of variance. Confirming the presence of an AB, the analysis indicated a significant effect of Lag,  $\underline{F}(3, 45) = 191.70$ , p < .001, MS<sub>e</sub> = 96.87,  $\eta^2 = .93$ . In addition, confirming the impression that AB magnitude varied with set-size, there was a significant main effect of Set-size,  $\underline{F}(2, 30) = 104.02$ , p < .001, MS<sub>e</sub> = 271.17,  $\eta^2 = .87$ , and a significant interaction between Lag and Set-size,  $\underline{F}(6, 90) = 31.69$ , p < .001, MS<sub>e</sub> = 75.81,  $\eta^2 = .68$ .

# Discussion

The objective of Experiment 3 was to verify that the visual search task used in Experiment 2 modulated T1 processing time. Consistent with this proposition, T1 response times increased as a function of an increase in T1 set-size. The results also indicate a strong relationship between T1 set-size and AB magnitude. This outcome affirms earlier findings that when T1 is not masked, T1 processing time modulates AB magnitude.

Collectively, the results from Experiments 1-3 are entirely consistent with predictions of bottleneck models. However, bottleneck models make more than a general prediction that T1 processing time and AB magnitude should be related. These models also make specific predictions about the sequence of processing events underlying this relationship. The goal of the remaining experiments was to test these predictions. The first prediction, examined in Experiment 4, is that an increase in T1 processing time should increase the period of delay for T2 in Stage 1. The second prediction, examined in Experiment 5, is that while delayed in Stage 1, T2 should be vulnerable to masking.

#### **EXPERIMENT 4**

The purpose of Experiment 4 was to test the prediction of bottleneck models that an increase in T1 processing time should delay processing of T2. This prediction was evaluated by measuring T2 response times in an experimental paradigm similar to that employed in Experiments 2 and 3. On each trial, observers were presented with a visual-search task as T1, followed at a constant interval by a single letter as T2. Neither target was masked. Observers were instructed to classify T2 as a vowel or a consonant as quickly as possible. Then, at their leisure, they were told to report whether T1 was a "C" or a "G". According to

bottleneck models, as T1 set-size increases, T2 response times should also increase, reflecting the longer delay experienced by T2 before gaining access to Stage 2.

# Method

<u>Participants</u>. Twelve undergraduate students (10 female) at the University of British Columbia participated for course credit. All participants reported normal or corrected-tonormal vision. None had participated in the previous experiments.

<u>Apparatus and Stimuli</u>. The apparatus and stimuli were identical to those used in Experiment 2 with one exception. The second target always consisted of one of eight letters - either a vowel (A, E, I, O) or a consonant (V, T, L, P).

<u>Procedure</u>. Stimulus presentation was identical to that used in the condition without a mask after T1 in Experiment 2, with four exceptions. First, on each trial, T1-T2 lag was held constant at a duration of 300 ms. Second, T2 was always either a vowel or a consonant, selected at random with replacement from the set of eight potential targets. Although, vowels and consonants were constrained to appear equally often, no constraints were placed on the number of times that each particular target letter could be presented. Third, rather than identifying T2 on each trial, observers were asked to classify it as a vowel or a consonant as quickly as possible by pressing one-of-two appropriately marked keys on the keyboard. Fourth, because the second task was speeded, observers were told to make the vowel-consonant judgement first, and then to identify T1 at their leisure.

#### Results

Mean percentages of correct identifications of T1 were 94.3, 92.5 and 89.7 for setsizes 1, 2, and 5 respectively. These results were analyzed in a one-way analysis of variance, which revealed a significant effect of Set-size,  $\underline{F}(2, 22) = 4.27$ ,  $\underline{p} = .03$ ,  $MS_e = 15.55$ ,  $\eta^2 = .28$ , confirming that overall T1 accuracy decreased significantly as set-size increased.

Trials on which T2 errors were made were discarded from the response-time analysis. This amounted to a total of 8.9% of trials at set-size 1, 11.5% of trials at set-size 2, and 15.3% of trials at set-size 5. Response times from all other trials were screened using the outlier procedure described by Van Selst and Jolicoeur (1994). This resulted in the removal of a further 3.7% of RTs at set-size 1, 4.9% at set-size 2, and 3.9% at set-size 5. The remaining data were used to calculate mean RTs at each set-size. These means are illustrated in Figure 7. The results were analyzed in a one-way analysis of variance, which revealed a significant

## Insert Figure 7 about here

\_\_\_\_\_

effect of Set-size, <u>F</u>(2, 22) = 10.24, <u>p</u> = .001, MS<sub>e</sub> = 18046.82,  $\eta^2$  = .48, confirming that mean T2 response times increased significantly with an increase in T1 set-size.

## Discussion

According to bottleneck models, an increase in T1 processing time should delay T2 for a corresponding duration in Stage 1. Consistent with this prediction, T2 response times increased as a function of T1 set-size. Bottleneck models also predict that during this delay, T2 should be vulnerable to masking. Put differently, these models suggest that as T1 processing time increases, T2 should be vulnerable to masking for longer. This prediction was tested in Experiment 5.

### **EXPERIMENT 5**

The purpose of Experiment 5 was to evaluate the prediction of bottleneck models that T2 should be vulnerable to masking while T1 is being processed in Stage 2. To test this prediction, a paradigm similar to that in Experiment 4 was employed. Each trial began with the presentation of T1, which was a visual search task, followed after a fixed interval by T2, which was a letter. Observers were told to determine whether T1 was the letter "C" or "G" and then to identify the T2 letter. As in previous experiments, only T2 was masked.

To estimate the period for which T2 was vulnerable to masking, a procedure was used similar to that employed by Zuvic, Visser, and Di Lollo (2000). The ISI between T2 and a trailing mask was dynamically varied using a staircase procedure (PEST; Taylor & Creelman, 1967) to keep T2 accuracy at 80%. When accuracy dropped significantly below this level, the ISI between T2 and its mask was increased. This allowed more time for T2 to be processed, thereby increasing accuracy. In contrast, when performance rose significantly above 80%, the ISI between T2 and its mask was decreased. This reduced the amount of processing time available for T2, thereby decreasing accuracy. By the end of a block of trials, a "critical" ISI between T2 and its mask could be calculated, which estimated the interval required by T2 for 80% correct identification. This interval provided an estimate of the duration for which T2 was vulnerable to masking.

### <u>Method</u>

<u>Participants</u>. Twelve undergraduate students (10 female) at the University of British Columbia participated for course credit. All participants reported normal or corrected-tonormal vision. None had participated in the previous experiments. <u>Apparatus and Stimuli</u>. The apparatus and stimuli were identical to those used in Experiment 2.

<u>Procedure</u>. Stimulus presentation was identical to that used in the condition without a mask after T1 in Experiment 3, with two exceptions. First, on each trial, T1-T2 lag was held constant at a duration of 300 ms. Second, the ISI between T2 and its mask was dynamically varied using the PEST staircase procedure in order to maintain T2 accuracy at approximately 80%. Separate staircase procedures were used at each T1 set-size.

### Results

Mean percentages of correct identifications of T1 were 98.2, 95.6 and 87.7 at setsizes 1, 2, and 5, respectively. The results were analyzed in one-way analysis of variance, which revealed a significant effect of Set-size,  $\underline{F}(2, 22) = 13.76$ ,  $\underline{p} < .001$ , MS<sub>e</sub> = 25.84,  $\eta^2 =$ .56, confirming that overall T1 accuracy decreased significantly as set-size increased.

Estimates of "critical" T2-mask ISI were based exclusively on those trials in which T1 had been identified correctly. The "critical" ISI for each observer was obtained by calculating the mean T2-Mask ISI of the last 50 trials at each set-size. Mean "critical" ISIs as a function of set-size are illustrated in Figure 8. An inspection of this figure suggests that "critical"

Insert Figure 8 about here

\_\_\_\_\_

\_\_\_\_\_

ISI increased directly with increases in T1 set-size. Confirming this impression, a one-way analysis of variance revealed a significant effect of Set-size,  $\underline{F}(2, 22) = 28.95$ ,  $\underline{p} < .001$ , MS<sub>e</sub> = 3862.87,  $\eta^2 = .73$ .

## Discussion

The goal of Experiment 5 was to verify that an increase in T1 processing time led to an increase in the period for which T2 was vulnerable to masking. Consistent with this proposition, the "critical" T2-Mask ISI required for 80% correct identification increased as a function of T1 set-size. This outcome is also consistent with previous findings that suggest that a target's vulnerability to masking varies as a function of the availability of attentional resources (see Appendix A).

#### GENERAL DISCUSSION

According to bottleneck models, temporal processing occurs in two broadly sequential stages – an initial parallel stage in which potential targets are detected, followed by a second capacity-limited stage responsible for memory encoding as well as responseplanning and execution. The focus of the present work was on evaluating these models by testing their predictions about the effect of increasing T1 processing time on AB magnitude, T2 processing time, and the vulnerability of T2 to masking.

In Experiments 1 to 3, the prediction that increasing T1 processing time would increase the magnitude of the AB was tested. Consistent with that prediction, an increase in T1 processing time was related to an increase in the magnitude of the AB. This was true in Experiment 1, in which a size-discrimination task was used as T1, as well as in Experiments 2 and 3, in which a visual-search task was used as T1.

Also notable is that the relationship between T1 processing time and AB magnitude was only obtained when T1 was not masked. With a mask after T1, the magnitude of the AB did not vary with T1 processing time. This difference is explicable in terms of interruption masking of T1. Because it occurred at a constant T1-mask SOA, this masking effectively equated T1 processing times across conditions. In accordance with predictions of bottleneck models, this equated the magnitude of the AB.

In Experiments 4 and 5, two additional predictions of bottleneck models were tested and confirmed. In Experiment 4, evidence was obtained suggesting that increases in T1 processing time were related to delays in T2 identification. This is consistent with predictions of bottleneck models that the period of delay experienced by T2 in Stage 1 is determined by the time required to process T1. Finally, in Experiment 5, it was demonstrated that the duration of T1 processing was directly related to the vulnerability of T2 to masking by a trailing stimulus. This is consistent with the prediction of bottleneck models that while delayed in Stage 1, T2 is vulnerable to masking.

Considered collectively, the results from the present experiments are highly supportive of bottleneck models. One issue that needs to be addressed, however, is whether other models of the AB could also account for these findings. In the next section, interference models of the AB will be analyzed in order to determine whether they can account for the present results.

## Interference Models of the AB

To explain temporal processing in vision, interference models (e.g. Shapiro & Raymond, 1994) propose a series of processing stages similar to those suggested by Duncan

& Humphreys (1989) to explain visual search. In an initial stage, representations of stimuli from across the visual field are constructed in parallel. These representations are then compared to a target template that is set to pass only stimuli relevant to the current task. For example, in a typical AB experiment where the task is to identify two letters from amongst a stream of non-letter distractors, the template would be configured to pass only letter-like stimuli.

Items that match the target template enter a visual short-term memory (VSTM) store where they compete for a weighting that is assigned from a limited pool of resources. The amount of weighting that an item receives is determined by a number of factors such as goodness-of-fit to the target template, temporal proximity to a previous target, and order of entry into VSTM. The importance of this weighting lies in the fact that it determines the order in which items gain access to high-level processing. Items with a high weighting are processed quickly, while items with a low weighting remain in VSTM for a relatively long period of time. During this time, they are subject to decay which reduces the likelihood that they will be successfully identified.

Within this processing framework, the AB is explained as follows. At shorter lags, four items compete for weighting in VSTM: T1, T2, and their respective masks. Because T1 and its mask are the first to enter, they receive a large proportion of the available weighting. This leaves T2 with a comparatively low weighting. As a result, on many trials, T2 decays from VSTM before it can gain access to high-level processing. In contrast, at longer lags, only two items compete in VSTM: T2 and its mask. This is because by the time T2 is presented, T1 and its mask have already gained access to high-level processing and been

flushed from VSTM. Because T2 has very little competition in VSTM, it receives a high weighting that allows it to gain quick access to high-level processing.

For the present purposes, the critical question is whether interference models can account for the fact that T1 processing time modulates AB magnitude. It is apparent from an analysis of the model that the magnitude of the AB is determined entirely by the weighting that items receive in VSTM. Therefore, any factor that changes an item weighting will impact the magnitude of the AB. For example, if a factor increases the weighting for T1 in VSTM, this will produce a relatively greater impairment in T2 performance. This is because the total amount of weighting in VSTM is fixed and therefore increasing the weighting of one item necessarily robs weighting from other items.

Importantly, although a number of factors have been suggested to influence the weighting of items in VSTM, such as order-of-entry, there is no suggestion that processing time is one of these factors. Therefore, it must be concluded that as they currently stand, interference models cannot account for the present findings. This is not to say that they will never be able to do so, but rather that modifications to the model must be made. Such modifications may take the form of additional processing stages, or simply a clarification of the functioning of the current processing architecture.

The present results are not alone in indicating that interference models require some modification or clarification. For example, Giesbrecht and Di Lollo (1998; see also Brehaut et al., 1999) examined the role of T2 masking in the AB. They found that the presence of the AB was critically dependent on the presence of an interruption mask after T2. Without a mask, there was no AB. Interference models cannot readily explain this finding because,

according to these models, the T2 mask should receive relatively little weight in VSTM due to its temporal position in the AB sequence.

Similarly, Chun (1997) noted that in a standard RSVP paradigm the majority of T2 misidentifications consist of the mask presented after T2. This finding is inconsistent with interference models which argue that T2 intrusions should primarily consist of items in VSTM that have a higher weighting than T2 – namely, either T1 or its mask. Collectively, the findings of Giesbrecht and Di Lollo (1998) as well as those of Chun (1997) point to an important role for the mask after T2. However, the nature of this role remains to be accounted for in the context of the interference models.

#### The Role of the T1 Mask

The preceding section touched on the role of masking in the AB. As noted, a number of studies have suggested that a mask after T2 is necessary in order for an AB to be obtained. However, there have also been numerous findings that have indicated that the mask after T1 is equally critical to obtaining an AB. As noted earlier, in their seminal work, Raymond et al. (1992) observed that while an AB occurred reliably when an item followed T1 by 100 ms, the effect was much reduced when this "+1" item was omitted. Expanding upon this earlier finding, both Seiffert and Di Lollo (1997) and Grandison et al. (1997) systematically examined the role of the T1 mask. They found that many types of T1 masking yielded an AB. For example roughly equivalent AB deficits were obtained by Seiffert and Di Lollo (1997) using integration, interruption, or metacontrast masking of T1. However, if the mask after T1 was omitted, little or no AB deficit was obtained. These findings contrast sharply with the present results. In Experiments 2 and 3, with a visual search task as T1, large ABs were obtained both with and without a mask after T1. Similarly, in Experiment 1, although the magnitude of the AB was much reduced when the mask after T1 was omitted, it was still highly significant. Thus, the present results suggest that while the T1 mask can influence the magnitude of the AB, unlike the mask after T2, it is not <u>necessary</u> for an AB to occur.

One explanation for the discrepancy between the present results and previous findings, that is in keeping with the bottleneck framework, is that it is not the presence or absence of a mask after T1 that is important <u>per se</u>, but rather the effect of the mask on the duration of T1 processing. A relationship between T1 processing time and the magnitude of the AB has been demonstrated repeatedly in the preceding experiments. These experiments have also shown that when the duration of T1 processing is relatively brief, the magnitude of the AB is quite small. Witness for example, the high level of T2 accuracy when T1 consisted of a single letter presented in a random location as in Experiments 2 and 3.

Notably, the T1 display in these experiments was quite similar to those used by Grandison et al. (1997), Raymond et al. (1992), and Seiffert and Di Lollo (1997) in which a single letter at fixation was to be identified. This suggests the possibility that in these experiments, like Experiments 2 and 3, T1 processing time without a mask was so brief as to leave T2 performance almost completely unimpaired. It was only with the addition of mask that T1 processing time became sufficiently lengthy to lead to a measurable level of impairment in T2 performance. Viewed from this perspective, it can be seen that it is not the presence or absence of a T1 mask that determines whether an AB occurs or not. An AB can arise either with or without a mask after T1. Rather, what is important is the processing time required by T1. To the extent that a T1 mask increases T1 processing time to a point where it will impair identification of a temporally trailing T2, then the T1 mask is "necessary" to produce an AB. However, if T1 processing time is already relatively lengthy, as was the case with the larger set-sizes in Experiments 2 and 3, the presence of a mask after T1 will do little to determine whether an AB does or does not occur.

#### Application to other Studies

In the present work, it has been argued that failures to find a relationship between T1difficulty and the AB stems from the use of a pattern mask after T1. It seems that a similar argument could be advanced to account for the results of previous studies such as Ward et al. (1997) and McLaughlin et al. (2001). Indeed, in the case of the later, the authors repeatedly refer to their masking of T1 as "interruption masking" (pg. 175). However, although plausible, this argument has one potential problem. Earlier, it was noted that pattern masks could be divided into two broad types – integration and interruption – based on the way in which they interfere with target identification. This interference is conventionally thought to be a function of the SOA between the target and the mask. According to this "SOA law" (Kahneman, 1968), when target-mask SOA is relatively short, interference will primarily take the form of integration masking; on the other hand, when the SOA is relatively long, interruption will dominate. Given this division, an examination of the presentation conditions in McLaughlin et al. (2001) suggests that integration masking predominated in the "hard" (SOA = 30 ms), and "medium" (SOA = 45 ms) conditions, while interruption masking was a significant factor only in the "easy" (SOA = 60 ms) condition. If this were the case, than it might be expected that differences in AB magnitude should have occurred because little or no interruption masking would have occurred in the "hard" and "medium" conditions. This would allow T1 processing to continue uninterrupted, thereby producing differences in AB magnitude.

Before drawing this conclusion, it is important to note that while the target-mask SOA was different in each condition, the total time between the onset of the target and the offset of the mask was identical. Recent findings, reported by Enns and Di Lollo (1997) and Di Lollo, Enns, and Rensink (2000), suggest that these presentation conditions can lead to object substitution masking – a type of interference akin to masking by interruption. In their experiments, Enns and colleagues found that for unattended targets, identification accuracy diminished rapidly as the duration of the mask was increased, even though target-mask SOA remained constant. This finding suggests that it is not only the target-mask SOA that is important in determining the level of masking, but also the interval between the onset of the target and the offset of the mask. This raises the possibility that although integration masking may have occurred to varying degrees as a function of T1-difficulty in the experiments of McLaughlin et al. (2001), an equal degree of interruption masking arose due to the fact that the total duration of the target plus the mask remained constant.

An equally important question that arises from McLaughlin et al. (2001) concerns the role of attentional set in modulating the relationship between T1 processing time and AB

magnitude. In their experiment, no effect of T1-difficulty on AB magnitude was obtained when the level of T1-difficulty was randomly determined from trial-to-trial. Based on these findings, McLaughlin et al. (2001) speculated that T1-difficulty might modulate AB magnitude only when the same level of difficulty occurs consistently in a block of trials. According to McLaughlin et al. (2001), such blocking would allow observers to allocate more resources to T1 when it is a "hard" task, thereby robbing resources from T2. This would produce a difference in the magnitude of the AB as a function of T1-difficulty.

Results consistent with this prediction were obtained in the present work. In each of the experiments, the time required to process T1 was the same across a block of trials. However, to determine whether this pattern of results is actually attributable to differences in attentional set, they must be compared to results from an identical set of experiments in which T1 processing time was varied from trial to trial. If the effect of T1 processing time disappears when processing time is varied, thereby preventing an attentional set from being formed, it would provide strong support for the role of attentional set in modulating the relationship between T1 processing time and AB magnitude. In a number of recent studies (Visser, in preparation), Experiments 2 and 3 were replicated but with set-size varied randomly from trial-to-trial. The results were almost identical to those from Experiments 2 and 3. This outcome suggests that attentional set played, at best, a minimal role in determining the outcome of the present work.

### Generalizability of Bottleneck Models

The present results indicate that bottleneck models provide a parsimonious account of the AB. However, beyond the study of this specific phenomenon, the more general goal of

the present work is to uncover a plausible model of temporal processing in vision. Thus, what must be asked is not merely whether bottleneck models can account for a single phenomenon in vision such as the AB, but whether they might also provide some insight into other extant findings. A number of recent studies have suggested that at least two other phenomena, contingent capture and task-switching, may be explained within a bottleneck framework.

On a typical trial in a contingent capture experiment, observers are presented with a cue display, followed after a variable interval by a target display. For example, in the cue display, a circle might appear in one location on the screen, while in the target display, a letter to be identified will appear in either the same location as the cue or in a different location. Importantly, although the cue and the target share some dimension, such as colour, the location of the cue is completely non-predictive of the location of the target. However, although observers are told that the cue is non-informative, it still exerts a significant influence. When the cue and target are in different locations, target detection is significantly slower than when they are in the same location.

Conventionally, this slowing has been attributed entirely to shifts in spatial attention (e.g. Folk, Remington, & Johnston, 1992). It is said that because the cue shares features with the target, the onset of the cue compels a shift in the "spotlight" of attention to its location. Then, when the target appears in a different location, the spotlight must be moved again to the target location. Because two spatial shifts are necessary, target identification is slower than when the cue and target appear in the same location, and only a single shift to the cued location must be performed. Recent evidence, however, suggests that at least some of the slowing that occurs when cue and target appear in different locations may be due to processing of the cue. Zuvic, Visser, and Di Lollo (submitted) presented an RSVP stream of items in the center of a display, followed by a single target letter to be identified in the periphery. The RSVP stream could consist either of letters or patterns of random dots. The results were unequivocal. When the stream consisted of letters, identification of the peripheral target was much slower than when the stream consisted of random dots.

Zuvic et al. (submitted) interpreted this deficit in terms of inadvertent processing of the central letter distractors. They argued that when the distractors shared features with the target, as was the case when both were letters, the distractors were compellingly processed. As a result, when the target letter appeared, it was prevented from gaining immediate access to processing resources because these resources were busy processing an irrelevant distractor. This resulted in slower target response times. In contrast, when the distractors did not share features with the target, as was the case when distractors were random dots, distractors went largely unprocessed. As a result, when the target appeared, it gained immediate access to high-level processing, thus resulting in faster response times.

This explanation for contingent capture is clearly quite similar to that advanced to explain the AB. In essence, the source of impairment is the same in both deficits – a target item is prevented from gaining immediate access to a high-level stage of processing because these processing resources are occupied with another item. The primary difference is in the nature of that item. In the case of the AB, the item is a prior target. In the case of contingent capture, the item is a distractor that is compellingly processed because it shares target

features (for a similar effect due to processing of distractors in the AB see Chun, 1997; Visser, Bischof, & Di Lollo, submitted).

Another phenomenon that can be accounted for in terms of bottleneck models is taskswitching. Quite often, in order to perform activities, it is necessary to rapidly switch between different tasks within a short time period. To take a simple case, when playing the outfield in baseball, the ball must first be caught and then thrown back to the infield. This requires an initial period in which the baseball must be tracked in order for it to be caught, and then a rapid shift in order to locate the infielder and throw the ball to them. Recently, Visser, Bischof, and Di Lollo (1999) noted that such rapid task-switching may be studied in the context of a typical AB experiment by focusing on a phenomenon referred to as "lag-1 sparing".

In a number of published experiments on the AB, it has been found that when T2 directly follows T1 (i.e. the ordinal position referred to as "lag-1"), second-target accuracy is relatively high compared to when an intervening item separates T1 and T2. This pattern of results has been termed "lag-1 sparing" to denote the fact that T2 accuracy seems to be relatively "spared" compared to later lags. What is curious about lag-1 sparing, however, is that it only occurs in about half of the published studies. In the other half, performance at lag-1 is the poorest of all the lags.

In order to explain this discrepancy, Visser et al. (1999) suggested that the presence of lag-1 sparing indicates that observers can successfully switch between T1 and T2 while the absence of lag-1 sparing indicates they cannot. For example, if T1 is a digit to be identified, and T2 is a letter to be identified, a category switch is required between digits and letters. If this switch can be accomplished in the interval separating T1 and T2, then both targets will gain access to high-level processing. As a result, both targets are processed, and lag-1 sparing will occur. In contrast, if the switch cannot be accomplished, only T1 will gain access to high-level processing. This will force T2 to remain delayed at an earlier stage of processing where it is vulnerable to decay or masking. As a result, lag-1 sparing will not occur.

### A Re-entrant Approach to Temporal Processing

Although the bottleneck models provide a parsimonious account for the AB and other related phenomena, this account is only a descriptive one – it does not, for example, provide insight into the brain mechanisms that might underlie the AB. The goal of this section is to speculate about such brain mechanisms, and to understand how these mechanisms could account for the importance of T1 processing duration in determining the magnitude of the AB. The starting point for this speculation is the object-substitution framework (Enns & Di Lollo, 1997; Di Lollo et al., 2000) developed to explain visual masking.

The majority of theories of visual perception view the process as feed-forward. That is, visual perception is seen as progressing through a series of hierarchical stages. This begins with transduction of light into electrical impulses in the eye, continues through to the decomposition of the image into simple features in V1, and then ends with a re-assembly of these simple features into consciously-perceived objects in higher brain regions. Importantly, progress through the stages is viewed as unidirectional, with information being passed from lower stages to higher stages, but not vice-versa. With respect to theories of temporal processing in vision, the feedforward viewpoint has led to several vexing questions. Chief amongst these concerns is the mechanisms that underlie backward masking – that is, interference with target identification caused by a trailing visual mask. Put simply, the problem is understanding how the neural impulses corresponding to the mask can "catch up" to neural impulses of the target. Such "catch-up" would seem to be required in a feedforward model in order to explain how a trailing mask could interfere with a preceding target.

One solution to this problem, suggested by Breitmeyer and Ganz (1976; Breitmeyer, 1984), is a model based on the functioning of transient and sustained pathways in vision. These pathways are known to be anatomically and functionally separate. The transient pathway carries information about rapid onsets of stimuli and carries information necessary for motion perception. The sustained pathway carries information about stimuli that remain visible for a period of time, and thus supports object recognition. Critically, information is known to travel faster along the transient than along the sustained pathway. Breitmeyer and Ganz (1976) argued that backward masking arises when activation in the sustained pathway, carrying information necessary for target identification, is suppressed by activation carried along the transient pathway, signaling the onset of the mask. Such interference is possible because transient signals are transmitted faster than sustained signals, thus allowing a neural representation of the mask to "catch up" to the representation of the preceding target.

Although this explanation for backward masking does solve the problem of how a mask presented after a target can impair its identification, it cannot account for a more recent phenomenon, discussed earlier, known as common-onset masking (Enns & Di Lollo, 1997).

On a typical trial in this paradigm, a target is presented simultaneously with a mask. Importantly, the mask remains on the screen for a variable duration after the target is turned off. For example, in Enns and Di Lollo (1997), the target was a Landolt C, while the mask consisted of four dots surrounding the target, and forming the corners of an imaginary square. Two factors are varied systematically – the duration for which the mask remained on the display after the offset of the target, and the number of distractors presented along with the target. When the mask and target terminated simultaneously, or when there were no distractors, the four dots did not mask the target. However, as the number of distractors presented with the target increased, masking began to increase steeply as a function of the time for which the mask remained on the display after the offset of the target.

Two aspects of this result are difficult to account for if it is assumed that masking can be caused only by transient-on-sustained inhibition. First, being presented simultaneously with the target, the mask cannot generate a transient signal separate from the one triggered by the target. Rather, target and mask together generate a single transient. Thus, it is impossible for the mask transient to suppress the sustained activity caused by the target. Second, the fact that the magnitude of masking increases with the duration of the mask also argues against some sort of transient-on-sustained inhibition resulting from the offset of the target. If anything, inhibition from the offset of the mask should be at its maximum when target and mask terminate simultaneously. Rather than a new stimulus (i.e. the mask) interfering with a prior stimulus (i.e. the target), common-onset masking seems to represent a case in which parts of a single stimulus interfere with each other when some parts remain on view after others have disappeared. Another aspect of common-onset masking that is relevant to the present work is its dependence on attention. When the target and mask were presented alone, the mask could be left on indefinitely, without impairing target identification. It was only when the target was presented along with distractors, thereby preventing attentional resources from being focused on it, that masking occurred. A review of the masking literature (see Appendix A) suggests that much of the effectiveness of visual masking may be attributable to such attentional effects.

In order to explain the importance of attention in masking, and the fact that it could occur independently of transients, Di Lollo et al. (2000) proposed a theory of masking based on a fundamentally different view of visual processing. Rather than a hierarchical system of processing in which information is passed in only one direction, they suggested a visual processing system based on the principle of re-entrant processing in which information flows in two directions – upwards from lower to higher visual centers, and then downwards from higher to lower visual centers.

There is abundant evidence to show that re-entrant pathways in the visual system, and indeed in the brain in general, are ubiquitous (e.g. Felleman & Van Essen, 1991; Perkel, Bullier, & Kennedy, 1986; Shipp & Zeki, 1989; Sillito, Jones, Gerstein, & West, 1994; Zeki, 1993). However, it is only relatively recently that such re-entrant activity has been ascribed any importance in visual processing. For example, Lamme, Rodriguez, and Spekreijse (1999; Lamme & Spekreijse, 1999) recorded single-cell activity in a macaque monkey while it was viewing a single rectangle on a plain background. They found that individual cells in V1 initially responded only to line segments of a preferred orientation in their visual field.

However, within 80 ms of the onset of the stimulus, the same cells began to respond only if the preferred line segment formed the boundary of the figure. Finally, approximately 140 ms after the onset of the stimulus, the cells responded to segments of their preferred orientation only if they were part of the figure but not if they were part of the background. Lamme et al. (1999) interpreted these results to indicate that the function of cells in V1 was being actively modulated by re-entrant signals from higher visual areas such as V4. This modulation altered the sensitivity of the cells in V1 such that their function changed from a line detector, to an edge detector, and finally to a figure detector in a short span of time after the initial presentation of the visual target.

Applying a similar logic to visual masking, Di Lollo et al. (2000) suggested that processing of a target stimulus begins with an initial decomposition into simple features at low levels in the visual system (e.g. V1). This information then cascades forward to higher areas in the visual system that generate "perceptual hypotheses" about the identity of the target. For example, given a target letter "T", potential perceptual hypotheses about the target's identity may include not only "T' but also stimuli with similar features such as "L" or "7". These perceptual hypotheses are correlated with available information about stimuli in lower visual areas via a process of iterative correlation that continues until a critical correlation is achieved. This correlation represents the confirmation of a perceptual hypothesis and results in conscious identification of the target.

According to this re-entrant model, masking occurs when a representation of the mask overwrites the representation of the target in lower visual areas before the perceptual hypothesis about the target's identity has been confirmed. As a result of this overwriting,

referred to as masking by <u>object substitution</u>, a new perceptual hypothesis is created that corresponds to the identity of the mask alone. This results in conscious perception of the mask. It is important to note, however, that not all information about the target is obliterated. Evidence from subliminal priming literature (e.g. Marcel, 1983; Visser, Merikle, & Di Lollo, 1998) suggests that some lingering effects of target processing can influence perception of subsequent stimuli.

This theory provides an account for common-onset masking if it is assumed that the perceptual hypothesis about the target surrounded by the four dots cannot be confirmed before the target is removed from the display. Under these circumstances, the target representation begins to rapidly decay in V1, while the representation of the four dots remaining on the screen is unaffected. This causes a rapid decrease in the correlation between the perceptual hypothesis corresponding to the target, and the information available in lower visual areas about what is actually being perceived. Ultimately, if the dots remain on the screen for a sufficient duration after the offset of the target, a new perceptual hypothesis is created corresponding to the four dots alone. This leads to masking of the target stimulus, as only the four dots are consciously perceived.

Although this theory provides an excellent account of common-onset masking, the more important question for the present work is whether it can also explain the AB. As noted earlier, the presence of the AB depends critically on masking of T2. Without a mask, T2 is identified accurately even when presented quite closely in time to a prior target. This finding suggests that the AB arises from the synergistic action of two factors: inattention and backward masking (see Visser, Merikle, & Di Lollo, in press). Given that these are the same

two factors already shown to be necessary to produce object substitution masking, the inference is almost compelling that the AB can be explained by a process akin to object substitution.

One plausible scenario is that while a re-entrant loop is ongoing for T1, no other reentrant processing can occur for other incoming stimuli (Luck, 2000). As a result, when T2 arrives soon after T1, its representation must wait in lower visual areas for re-entrant processing of T1 to be completed. While waiting, this representation of T2 is vulnerable to being overwritten by the trailing mask. As a result, by the time processing of T1 has been completed, there may no longer be an identifiable representation of T2 available, and T2 will fail to be identified.

Such an explanation for the AB on the basis of re-entrant processing provides an elegant account of the basic phenomenon. It also accounts for a variety of empirical findings in the AB literature. For example, it was noted earlier that the majority of T2 errors consist of intrusions by the item directly after T2 (Chun, 1997). This is easily explainable within the re-entrant processing model since it is the item directly after the unattended target that is thought to replace it in lower visual areas, and ultimately in conscious awareness. Similarly, the re-entrant processing model can also explain why T1 is never masked, while T2 masking varies as a function of temporal lag. Because it is the first target, processing resources are maximally available for T1. This means that the tasks of forming perceptual hypotheses and iterative correlations can be performed before T1 is overwritten by its mask. In contrast, because processing resources are occupied with T1 when T2 is presented at shorter lags, object substitution can occur.

Lastly, this framework also provides a natural account of the influence of T1 processing time. If it is assumed that at least part of the increase in processing time is due to an increase in the duration of re-entrant processing for T1, this would translate into a longer delay for the representation of T2 in lower visual areas before it can be re-iteratively processed. As noted earlier, this would make T2 vulnerable to overwriting by the trailing mask for longer, and decrease the probability of its identification.

In summary, one plausible means of instantiating bottleneck models of the AB in terms of brain mechanisms comes from a re-entrant processing model such as that of Di Lollo et al. (2000). However, at the present time, much more research is required before a detailed theory can be created and verified empirically. Chief amongst the questions to be answered is the exact means by which processing of T1 interferes with processing of T2.

# References

- Brehaut, J., Enns, J. T., & Di Lollo, V. (1999). Visual masking plays two roles in the attentional blink. <u>Perception and Psychophysics</u>, <u>61</u>, 1436-1448.
- Breitmeyer, B.G., & Ganz, L. (1976). Implications of sustained and transient channels for theories of visual pattern masking, saccadic suppression, and information processing.
  <u>Psychological Review</u>, 83, 1-36.
- Breitmeyer, B. G. (1984). <u>Visual masking: An integrative approach</u>. New York: Oxford University Press.
- Broadbent, D. E., & Broadbent, M. H. (1987). From detection to identification: Response to multiple targets in rapid serial visual presentation. <u>Perception & Psychophysics</u>, <u>42</u>, 105-113.
- Chun, M. M., & Potter, M. C. (1995). A two-stage model for multiple target detection in rapid serial visual presentation. <u>Journal of Experimental Psychology: Human</u> <u>Perception and Performance, 21</u>, 109-127.
- Chun, M. M. (1997). Temporal binding errors are redistributed in the attentional blink. <u>Perception & Psychophysics</u>, <u>59</u>, 1191-1199.
- Di Lollo, V., Enns, J. T., & Rensink, R. (in press). Psychophysics of re-entrant pathways: A computational model of object substitution (CMOS). <u>Journal of Experimental</u> <u>Psychology: General</u>.
- Duncan, J., Martens, S., & Ward, R. (1997, June 19). Restricted attentional capacity within but not within sensory modalities. <u>Nature</u>, <u>387</u>, 808-810.

- Duncan, J. & Humphreys, G. (1989). Visual search and stimulus similarity. <u>Psychological</u> <u>Review, 96</u>, 433-458.
- Enns, J. T., & Di Lollo, V. (1997). Object substitution: A new form of masking in unattended visual locations. <u>Psychological Science</u>, 8, 135-139.
- Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in primate visual cortex. <u>Cerebral Cortex</u>, 1, 1-47.
- Folk, C. L., Remington, R. W., & Johnston, J. C. (1992). Involuntary covert orienting is contingent on attentional control settings. <u>Journal of Experimental Psychology:</u> <u>Human Perception and Performance, 18</u>, 1030-1044.
- Giesbrecht, B. L., & Di Lollo, V. (1998). Beyond the attentional blink: Visual masking by object substitution. Journal of Experimental Psychology: Human Perception and <u>Performance</u>, 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. <u>Perception & Psychophysics</u>, <u>59</u>, 266-274.
- Jolicoeur, P. (1998). Modulation of the attentional blink by on-line response selection: Evidence from speeded and unspeeded Task-sub-1 decisions. <u>Memory and</u> <u>Cognition, 26</u>, 1014-1032.
- Joseph, J. S., Chun, M. M., & Nakayama, K. (1997, June 19). Attentional requirements in a preattentive feature search task. <u>Nature</u>, <u>387</u>, 805-808.
- Kahneman, D. (1968). Method, findings, and theory in studies of visual masking. <u>Psychological Bulletin</u>, 70, 404-425.

- Lamme, V. A. F., Rodriguez, V. A., & Spekreijse, H. (1999). Separate processing dynamics for texture elements, boundaries and surfaces in primary visual cortex of the macaque monkey. Cerebral Cortex, 9, 406-413.
- Lamme, V. A. F., & Spekreijse, H. (1999). Contextual modulation in V1 and scene perception. In M. Gazzaniga (Ed.) <u>The Cognitive Neurosciences</u>. Cambridge: MIT Press.
- Logan, G. D. (1994). Spatial attention and the apprehension of spatial relations. <u>Journal of</u> <u>Experimental Psychology: Human Perception and Performance, 20</u>, 1015-1036.
- Luck, S. J., Vogel, E. K., & Shapiro, K. L. (1996). Word meanings can be accessed but not reported during the attentional blink. <u>Science</u>, <u>383</u>, 616-618.
- Luck, S. J. (2000). Electrophysiological perspectives on the binding problem. Paper presented at the 41st Annual Meeting of the Psychonomic Society, New Orleans, LA.
- Maki, W. S., Frigen, K., & Paulson, K. (1997). Associative priming by targets and distractors during rapid serial visual presentation: Does word meaning survive the attentional blink? Journal of Experimental Psychology: Human Perception and Performance, 23, 1014-1034.
- Marcel, A. J. (1983). Conscious and unconscious perception: Experiments on visual masking and word recognition. <u>Cognitive Psychology</u>, <u>15</u>, 197-237.
- McLaughlin, E. N., Shore, D. I., & Klein, R. M. (2001). The attentional blink is immune to masking-induced data limits. <u>Quarterly Journal of Experimental Psychology</u>, <u>54A</u>, 169-196.

- Michaels, C. F., & Turvey, M. T. (1979). Central sources of visual masking: Indexing structures supporting see at a single, brief glance. <u>Psychological Research</u>, <u>41</u>, 1-61.
- Nazir, T. A. (1992). Effects of lateral masking and spatial precueing on gap-resolution in central and peripheral vision. <u>Vision Research</u>, 32, 771-777.
- Perkel, D. J., Bullier, J., & Kennedy, H. (1986). Topography of the afferent connectivity of area 17 in the Macaque monkey. Journal of Comparative Neurology, 253, 347-402.
- Ramachandran, V. S., & Cobb, S. (1995). Visual attention modulates metacontrast masking. Nature, 373, 66-68.
- Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink? <u>Journal of Experimental</u> Psychology: Human Perception and Performance, 18, 849-860.
- Ross, N.E., & Jolicoeur, P. (2000). Attentional blink for color. <u>Journal of Experimental</u> <u>Psychology: Human Perception and Performance, 25</u>, 1483-1494.
- 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., & Raymond, J. E. (1994). Temporal allocation of visual attention: Inhibition or interference? In D. Dagenbach & T. H. Carr (Eds.) <u>Inhibitory Mechanisms in</u> <u>attention, memory and language</u> (pp. 151-187). San Diego, CA: Academic Press.
- Shapiro, K. L., Raymond, J. E., & Arnell, K. M. (1994). Attention to visual pattern information produces the attentional blink in RSVP. <u>Journal of Experimental</u> <u>Psychology: Human Perception and Performance</u>, 20, 357-371.

- Shipp, S., & Zeki, S. (1989). The organization of connections between areas V5 and V1 in the macaque monkey visual cortex. European Journal of Neuroscience, 1, 309-332.
- Sillito, A. M., Jones, H. E., Gerstein, G. L., & West, D. C. (1994). Feature-linked synchronization of thalamic relay cell firing induced by feedback from visual cortex. <u>Nature, 369</u>, 479-482.
- Spencer, T. J., & Shuntich, R. (1970). Evidence for an interruption theory of backward masking. <u>Journal of Experimental Psychology</u>, <u>85</u>, 198-203.
- Taylor, M. M. & Creelman, C. D. (1967). PEST: Efficient estimates on probability functions. <u>The Journal of the Acoustical Society of America</u>, <u>41</u>, 782-787.
- Treisman, A., & Gelade, G. (1980). A feature integration theory of attention. <u>Cognitive</u> <u>Psychology</u>, <u>12</u>, 97-136.
- Turvey, M. T. (1973). On peripheral and central processes in vision: inferences from an information-processing analysis of masking with patterned stimuli. <u>Psychological</u> <u>Review</u>, <u>80</u>, 1-52.
- Van Selst, M., & Jolicoeur, P. (1994). A solution to the effect of sample size on outlier elimination. <u>The Quarterly Journal of Experimental Psychology</u>, <u>47A</u>, 631-650.
- Visser, T. A. W., Bischof, W. F., & Di Lollo, V. (1999). Attentional switching in spatial and non-spatial domains: evidence from the attentional blink. <u>Psychological Bulletin</u>, <u>125</u>, 458-469.
- Visser, T. A. W., Bischof, W. F., & Di Lollo, V. (submitted). Rapid-serial visual distraction: Task-irrelevant items can produce an attentional blink. <u>Perception and</u> Psychophysics.

- Visser, T. A. W., Merikle, P. M., & Di Lollo, V. (1998, November). A tale of two masks: Linking unconscious perception, the attentional blink, and masked priming. Paper presented at the 39th Annual Meeting of the Psychonomic Society, Dallas, TX.
- Ward, R., Duncan, J., '& Shapiro, K. (1997). Effects of similarity, difficulty, and nontarget presentation on the time course of visual attention. <u>Perception & Psychophysics</u>, <u>59</u>, 593-600.
- Zeki, S. (1993). A vision of the brain. Oxford: Blackwell.
- Zuvic, S. M., Visser, T. A. W., & Di Lollo, V. (2000). Direct estimates of processing delay in the attentional blink. <u>Psychological Research/Psychologische Forschung</u>, <u>63</u>, 192-198.
- Zuvic, S. M., Visser, T. A. W., & Di Lollo, V. (submitted). Focal distraction: Spatial shifts of attention are not required for contingent capture. <u>Journal of Experimental</u> <u>Psychology: Human Perception and Performance.</u>



<u>Figure 1.</u> Mean response times for target discrimination in the "easy" and "hard" conditions in Experiment 1A. Error bars represent one standard error of the mean.



Figure 2. Twenty-six "pseudo-letter" figures used as masks in Experiment 1B.



<u>Figure 3.</u> Mean accuracy of T2 identification, given correct identification of T1, as a function of the temporal lag between T1 and T2. Closed circles represent scores when the first task was "short". Closed triangles represent scores when the first task was "long". Error bars represent one standard error of the mean. Solid lines represent the results when T1 was not masked. Dashed lines represent results from the identical conditions with a mask after T1.


Figure 4. Mean accuracy of T2 identification, given correct identification of T1, as a function of the temporal lag between T1 and T2. Closed circles, triangles, and squares represent set-sizes 1, 2, and 5 respectively. Error bars represent one standard error of the mean. Panel A depicts T2 accuracy with a mask after T1. Panel B depicts T2 accuracy without a mask after T1.



Figure 5. Mean response times to identify T1, given correct identification of T1, as a function of the set-size of the T1 display. Error bars represent one standard error of the mean.



Figure 6. Mean accuracy of T2 identification, given correct identification of T1, as a function of the temporal lag between T1 and T2. Closed circles, triangles, and squares represent scores at T1 set-sizes of 1, 2, and 5 respectively. Error bars represent one standard error of the mean.



Figure 7. Mean response times to identify T2, given correct identification of T1 and T2, as a function of the set-size of the T1 display. Error bars represent one standard error of the mean.



Figure 8. Mean ISI between T2 and its mask required for 80% correct T2 letter identification, as a function of the set-size of the T1 display. Error bars represent one standard error of the mean.

# Appendix A: The Influence of Attention on Visual Masking

Our visual world is filled with objects that frequently appear in close spatial and temporal proximity. In the majority of cases, this proximity does not impair our ability to perceive objects in everyday viewing. Witness, for example, our ability to drive a motor vehicle in traffic or hit a baseball. Nonetheless, under some conditions, a phenomenon called <u>visual masking</u> arises in which the perception of an object is impaired by other objects in close spatial or temporal proximity. One example of masking can be seen when two objects are presented in rapid succession in the same spatial location. Under these circumstances, perception of the initial object, called the target, is impaired by the presentation of the trailing object, called the mask.

A great deal of research, spanning over 100 years, has sought to determine the conditions under which masking can occur. The primary focus of this research has been an examination of relatively low-level factors. These factors have included such things as target and mask duration, stimulus-onset asynchrony (SOA) between mask and target, as well as luminance and spatial frequencies of targets and masks (see Breitmeyer, 1984 for specific examples). However, much less work has been done to examine the potential influence of high-level factors, such as visual attention, on masking. In light of recent work that has shown that attention modulates a number of potentially related visual phenomena such as brightness perception (e.g. Enns, Brehaut, & Shore, 1999), spatial resolution (e.g. Yeshurun & Carrasco, 1999), and temporal integration (e.g. Visser & Enns, submitted), examination of the influence of high-level factors assumes considerable importance.

The few studies that have examined this issue strongly suggest that attention can modulate the magnitude of masking. One representative study was conducted by Di Lollo, Enns, and Rensink (in press; see also Brehaut, Enns, & Di Lollo, 1999; Enns & Di Lollo, 1997; Giesbrecht & Di Lollo, 1998; Visser, Zuvic, Bischof, & Di Lollo, 1999). In the experiment of Di Lollo et al. (in press), observers were shown a target surrounded by four dots. After approximately 10 ms, the target disappeared, leaving only the dots on display for a variable duration. When the target was the only item in the display, the presence of the dots did little to impair its identification, regardless of how long the dots remained in view. However, when other items were presented along with the target, identification accuracy declined sharply as a function of the number of accompanying items and the duration for which the dots remained on view. This pattern of results strongly suggests that the magnitude of masking was modulated by the availability of attention for the target. Namely, when presented alone, attention could be focused rapidly on the target allowing it to escape masking. In contrast, when the target was presented amongst distractors, attention could not be focused on the target as rapidly, leaving it vulnerable to masking.

The finding that attention modulates masking has important implications for theories of visual masking, and visual perception in general. Current theories of masking focus primarily on explaining the influence of low-level factors on target identification. For example, the transient-sustained theory of Breitmeyer and colleagues (e.g. Breitmeyer, 1984; Breitmeyer & Ganz, 1976) attributes masking to an inhibition of the target's sustained neural activity by transient neural activity generated by the onset of the mask. This provides an explanation for the influence of a temporally-trailing mask on target identification. Other theories focus on the importance of contour integration between the target and mask (e.g. Eriksen, 1966; Turvey, 1973). These theories explain the importance of factors such as stimulus luminance and spatial frequency. Importantly, however, none of these accounts provides for a role of attention on masking. This suggests that such theories either require revision or abandonment in favour of new accounts.

Beyond its relevance for theories of visual masking, the relationship between attention and masking implies a potentially fundamental role for masking in visual processing. Given that masking seems to occur only for unattended objects, one implication is that masking is a process by which old, and unwanted visual information (i.e. unattended stimuli) is overwritten by new, potentially relevant information (i.e. the mask). Without such a process, processing would be plagued by interference as the decaying representations of old stimuli competed for processing resources with representations of new stimuli (see Averbach & Sperling, 1961). Such a system would not only be inefficient but would lead to "smearing" of stimuli presented in rapid succession.

Although these suggestions are intriguing, before examining more carefully the role of attention and masking in visual perception, a more detailed understanding of the interplay between these two factors is required. Masking is an exceedingly complex phenomenon with multiple causes and multiple loci within the visual system. Given this complexity, the current number of studies is simply insufficient to allow a complete analysis of the interplay between attention and masking. However, much additional information may be gained by examining published papers that have used paradigms that are "attentional" in nature. Such studies were not designed specifically to examine the role of attentional factors, but used experimental manipulations that can be interpreted as varying the degree of attention deployed to masked stimuli. The goal of the present work is to examine these studies systematically in order to clarify the influence of attention on three different types of masking (reviewed below), and to integrate this information into existing theoretical frameworks of visual perception and masking.

In Section 1, I review three different types of visual masking. Next, in Section 2, I review a number of different types of attentional manipulations that have been used in visual masking experiments. Together, the different types of masking and attentional manipulations form a taxonomic scheme that will be used to organize and classify the empirical literature. This taxonomic scheme is outlined in Section 3, and applied to the literature in Section 4. Finally, in Section 5, I discuss a number of issues raised by an analysis of the literature, as well as ways of integrating these findings into a coherent theory of masking based upon the principle of object substitution outlined by Di Lollo, Enns, and Rensink (in press).

## **SECTION I: Types of Visual Masking**

As noted earlier, masking is a complex phenomenon with multiple causes. It can be produced by a number of different types of visual stimuli, and can occur at many different points in the stream of visual processing. One example of this diversity can be seen in the distinction between masking of a target pattern by a uniform light flash, and masking by a patterned stimulus. In the case of the light flash, degradation of the target occurs early in the stream of visual processing, likely before the optic chiasm (see Schiller, 1965; Smith & Schiller, 1966; Braddick, 1973). The masking effect has been attributed primarily to a reduction in the relative contrast of the target (Eriksen, 1966). In the case of the pattern mask, the mask can interact with the target representation at multiple points along the stream of visual processing, thereby impairing target identification in a number of different ways.

In the present work, I focus on three different types of pattern masking: integration masks, interruption masks, and lateral masks. These different types of masking can be functionally distinguished based on a number of characteristics including the spatiotemporal conditions that produce each type of masking, the underlying mechanisms responsible for masking, and the phenomenological percept that the mask produces.

#### **Integration Masking**

This type of masking occurs when a mask is presented in the same location as a target, and either precedes (forward masking) or follows (backward masking) the target by a stimulus-onset asynchrony of not-greater-than 100 ms. The function relating target-identification accuracy to target-mask SOA is U-shaped, with maximal masking at an SOA of 0 ms (i.e. when the target and mask are simultaneous). Given these temporal characteristics, integration masking has commonly been ascribed to an amalgamation of the contours of the target and the mask at a relatively early level in the visual system (Turvey, 1973). This amalgamation is akin to adding noise (the mask) to the signal (the target). Phenomenologically, the target appears to be present for a sufficient duration to be identified, but accurate identification is prevented by the presence of the extraneous mask contours (Liss, 1968).

## **Interruption Masking**

Unlike integration masking, interruption masking occurs only when a mask is presented after the target (backward masking), with maximum masking occurring at a targetmask SOA greater than 0 ms. Additionally, unlike integration masking, interruption masking can occur when the contours of the mask surround, but do not overlap the target. This latter type of masking is referred to as metacontrast (Alpern, 1953; Kahneman, 1967; Weisstein, 1972). Given these spatiotemporal characteristics, interruption masking has often been attributed to a halt in target processing caused by the presentation of the mask. This prevents complete processing of the target, thereby impairing its identification. Kolers (1968) likened this situation to that of a clerk (processing resources) faced with a series of customers who arrive aperiodically. When single customers arrive at infrequent intervals, the clerk can devote a substantial amount of time to each of them. However, when customers begin to enter rapidly, the time that the clerk can spend with each of them is correspondingly decreased. Just as in the store where a decrease in time results in less satisfied customers, in the visual system, a decrease in available processing time results in fewer target identifications. This theoretical perspective is bolstered by the phenomenological appearance of targets masked by interruption. Under these conditions, observers typically report that a clear target stimulus was presented but that it was not visible long enough to be successfully identified (Liss, 1968; Spencer, Hawkes, & Mattson, 1972).

## Lateral Masking

This type of masking possesses broad similarities with both integration and interruption masking. Lateral masking occurs only when the target and mask are presented

simultaneously in close spatial proximity. This produces a type of crowding of the contours of the target and mask that reduces the visibility of the target. The mechanisms responsible for lateral masking are unclear (see Breitmeyer, 1984; Nazir, 1992). However, the evidence suggests that both low- and high-level mechanisms are partially responsible. For example, there is abundant evidence that the strength of lateral masking increases as a function of the retinal eccentricity of the target and mask. This suggests that spatial resolution in the retina is of critical importance to the masking effect (e.g. Bouma, 1970). However, there is also evidence that interference between the target and mask can occur at the level of response planning. A number of experiments on the <u>flanker effect</u> conducted by Eriksen and colleagues (e.g. Eriksen, O'Hara, & Eriksen, 1982; Eriksen & Yeh, 1985) indicate that when targets and distractors that are associated with different motor responses appear within 1° of visual angle, interference begins to take place.

#### Summary

Pattern masking can be classified into three broad categories: integration masking, interruption masking, and lateral masking. Each of these masking types can be distinguished along multiple dimensions, and each seems to arise from different mechanisms. In the next section, I outline some common methodologies that have been used to vary attention in conjunction with these different types of masks. By factorially combining the different types of pattern masks and attentional manipulations, a taxonomic scheme can be devised which will be used as a basis for organizing and summarizing the existing masking literature.

### **SECTION II:** Types of attentional manipulations

A variety of methods have been used to vary the availability of attentional resources for processing a target stimulus. For the purposes of the present work, I will divide these manipulations into three broad categories: spatial, temporal, and non-spatiotemporal. Spatial manipulations of attention

Two different methodologies have been used to vary the spatial allocation of attention. In the case of <u>visual search</u> tasks, a display is presented that consists of a target and a variable number of distractors. This search display is followed by a masking display consisting of one or more masks, covering either the target alone or both the target and distractors. Under these circumstances, it is assumed that to identify the target, a serial search must be conducted, with the "spotlight" of attention being allocated to each stimulus in turn, until the target is found (e.g. Treisman & Gelade, 1980). Given this assumption, the more distractors added to the display, the lower the probability that attention will be allocated to the target location before the arrival of the masking display.

A second type of spatial manipulation is referred to as <u>location-cuing</u>. As in the visual search paradigm, a typical location-cueing experiment consists of a search display followed by a masking display. However, rather than varying the number of distractors to manipulate attentional availability for a target, what is varied instead is the validity of a cue presented prior to the search display. In the case of a valid cue, a marker appears that correctly indicates the location of the target in the forthcoming search display. This presumably allows the attentional "spotlight" to be efficiently deployed to the location of the target before its presentation. In the case of an invalid cue, the marker signals the location of

a distractor item instead of the target. This presumably causes observers to initially allocate their attentional focus to the wrong location, and then to have to reallocate attention to the correct location before the onset of the masking display. This reallocation process will reduce the likelihood of attention being deployed to the target. By comparing target identification in the valid- and invalid-cue conditions, the influence of attention on performance can be estimated.

### Temporal manipulations of attention

The availability of attention for target processing can be varied not only across space, but also over time. This latter type of manipulation is exemplified by the <u>attentional blink</u> paradigm, in which two targets are presented in rapid succession. Under these circumstances, identification of the first target is nearly perfect. However, identification of the second target varies as a function of inter-target lag, with performance being poorest at shorter lags, and gradually improving as lag is increased. This attentional blink deficit has commonly been ascribed to a lack of attentional resources for the second target, brought about by the requirement to attend to the first target (e.g. Chun & Potter, 1995; Jolicoeur, 1998; Raymond et al., 1992; Shapiro et al., 1994). On this account, the impairment is greatest at shorter lags because processing resources are most likely to be occupied with the first target. As lag increases, the magnitude of impairment decreases because processing of the first target is more likely to have been completed by the time the second target is presented.

## Non-spatiotemporal manipulations of attention

Three different types of non-spatiotemporal manipulations have been used in masking experiments. The first type, grouping manipulations, involve instructing observers to mentally group certain stimuli together. A good example of this methodology can be seen in Ramachandran and Cobb (1995). In one condition, observers were instructed to mentally group the target and masking stimuli together. In a second condition, observers were instructed to group the mask with an unrelated stimulus that was presented several degrees away from the target. Masking was much reduced in the second condition compared to the first – a result which Ramachandran and Cobb attributed to greater availability of attention for the target in the second condition. They argued that grouping the masks with another stimulus allowed attention to be allocated to the target more efficiently.

Another type of non-spatiotemporal manipulation of attention is <u>stimulus salience</u> in which a stimulus can be either highly meaningful or neutral. One example of this type of manipulation can be seen in the work of Mack and colleagues (e.g. Mack & Rock, 1998; Shelley-Tremblay & Mack, 2000). In one condition, the target word consisted of a highly-meaningful stimulus – the observer's own name. In a second condition, the target word was either a scrambled version of the observer's name or a regular word of similar length. The authors found that masking was much reduced when the target word consisted of the observer's name. They attributed this difference to the ability of the meaningful target to capture attention, thereby allowing it to be processed more efficiently, and escape masking.

A third type of non-spatiotemporal manipulation of attention is <u>target-distractor</u> <u>similarity</u>. It is assumed that the similarity between targets and masks influences the level of competition for common processing resources, including high-level resources such as visual attention. Highly similar targets and masks will presumably compete for resources more than less-similar targets. One example is Jacobson (1974) who used target shapes in conjunction with masks that were random dots, lines, or shapes. Masking was greatest when masks consisted of shapes, was reduced when masks were lines, and was nearly eliminated when masks were random dots. This suggests that when targets and masks were similar, competition for attentional resources reduced availability of attention for targets compared to the condition in which targets and masks were dissimilar.

# **SECTION III: Taxonomic Scheme**

By factorially combining the three types of pattern masking (integration, interruption, lateral) described in Section I with the three types of attentional manipulations (spatial, temporal, non-spatiotemporal) described in Section II, a 3x3 matrix is created that can be used to categorize the masking literature. A schematic of this matrix can be seen in Table A1. The

Insert Table A1 about here

numbers in each box indicate the number of studies that have examined the relationship between that particular type of masking and attentional manipulation.

In Section IV, I discuss the each type of masking (i.e. each column in Table A1) separately, with the discussion subdivided based on the different types of attentional manipulation.

## SECTION IV: Review of the literature on attention and masking

In total, the findings from 44 different papers are considered in this survey. Two aspects of this collection deserve special comment. First, it should be noted that the survey does not include a number of relevant articles that made use of the attentional blink paradigm. All of these papers demonstrate a relationship between attention and interruption masking, with the severity of masking increasing as attention is less available for secondtarget processing. However, in most cases, this result is treated as an incidental aspect of the results. For this reason, I chose to include only papers that specifically addressed the relationship between attention and masking in the context of the attentional blink paradigm.

The second comment concerns the distribution of papers within the taxonomic scheme. Inspection of Table A1 shows that the vast majority of studies in the survey are concerned with the relationship between attention and interruption masking. Clearly, this imbalance speaks to the need for more empirical work to be done. More importantly for the present purposes, however, the relatively small number of papers that have examined integration and lateral masking may limit the generality of the conclusions that can be drawn from the survey. Indeed this may be one factor that accounts for some of the mixed results that are reported below. Nonetheless, even in limited numbers, these studies indicate important aspects of the relationship between attention and masking that deserve further consideration.

### Interruption Masking and Spatial Manipulations of Attention

A summary of the studies that have examined this relationship is shown in Table A2. At first glance, it is readily apparent that these studies show a strong relationship between

Insert Table A2 about here

attention and masking. In each case, unattended targets are more vulnerable to masking. Two representative examples are the location-cuing studies of Krose and Julesz (1989) and the visual search experiments of Spencer and Shuntich (1970). In the former, Krose and Julesz (1989; Expt. 1) presented a circular array consisting of a single target and eleven distractors. Each item was then followed by a pattern mask at an SOA of 100 ms. When the location of the target was validly cued before the onset of the array, target identification was 20% better than when the target location was invalidly cued, indicating that attending to the target location helped it to escape masking. In the case of Spencer and Shuntich (1970), a target was presented either alone or with eleven accompanying distractors. A pattern mask followed each item in the display. Target identification was much better when the target was presented alone than when it was accompanied by distractors. Like the findings of Krose and Julesz, this result suggests that masking was reduced when attention could be efficiently allocated to the target.

With respect to spatial manipulations of attention, however, an alternate interpretation must also be considered. It is possible that the influence of distractors did not arise from the need to rapidly shift an attentional "spotlight" from location to location in a search for the

۰.

target. Rather, the addition of distractors to the display may have simply added decision noise by making it harder to decide which location contained the target, and which contained a distractor (see Cohn & Lasley, 1974; Palmer, Ames, & Lindsey, 1993). This does not require that attentional resources be any less available for target processing when distractors are present, but merely that targets and distractors are potentially confusable – as was the case in the visual experiments listed in Table 2. Although the present evidence cannot rule out this interpretation, the fact that other types of attentional manipulations, reviewed below, affect the magnitude of interruption masking suggests that an interpretation of these results in terms of attention is also warranted.

In addition to studies that have examined the influence of attentional availability for a target using the visual search paradigm, a number of studies have examined the influence of varying the number of masks presented after a search display. A representative example is a study conducted by Shiu & Pashler (1994) that presented a single target followed by either one or four masks. They found more masking with four masks than with one mask. As with the spatial manipulations described above, this finding has a number of interpretations. One possibility is that with a single mask, attention can be focussed on the mask, whereas with four masks, attention is divided amongst the masks. On this account, Shiu and Pashler's results suggest that as with unattended targets, unattended masks produce greater masking. However, this result can also be interpreted in terms of a decision-noise framework. With multiple masks, presented in close temporal contiguity to the targets, observers may have had some difficulty in determining the location of the target. Finally, a third possible interpretation is that by adding additional masks, the total amount of masking contours was

increased. This increase in the amount of masking contour may have impaired target identification to a greater extent. Once again, it is not possible, given the present results, to distinguish between these interpretations. However, as shown below, there is some evidence from studies using other types of attentional manipulations that availability of attention for the masks does influence the magnitude of interruption masking. Thus, there is some reason to believe that the effects of increasing the number of masks on the magnitude of interruption masking may be due to attentional factors.

#### Interruption Masking and Temporal Manipulations of Attention

These studies are listed in Table A3. As with studies that have used spatial manipulations of attention, the results using temporal manipulations of attention uniformly

-----·

Insert Table A3 about here

indicate that unattended targets are more vulnerable to interruption masking. A representative study is that of Giesbrecht and Di Lollo (1998) who compared the level of interruption masking for the fully-attended first target with the level of masking for the unattended second target in the attentional blink paradigm. For the first target, a pattern mask presented at an SOA of 100 ms produced little impairment. In contrast, for the unattended second target, an identical pattern mask presented at the same SOA led to a massive impairment in target identification. Moreover, the level of impairment for the second target varied with inter-target lag suggesting a direct link between attentional availability and the level of interruption masking.

Similar results have been obtained using a paradigm known as "unmasking". For example, Michaels and Turvey (1979; Expt. F1) presented three targets: T1, T2, and T3. The temporal separation between them was such that when only T1 and T2 were presented, T1 was masked by T2. When only T2 and T3 were presented, T2 was not masked by T3. Critically, however, when T1, T2, and T3 were presented in sequence, T2 was masked while T1 was not. This finding indicates that, as in the attentional blink paradigm, when T2 was not preceded by a stimulus, it was fully attended and thus escaped interruption masking by T3. However, when T2 was preceded by T1, attentional resources were less available for processing of T2. As a result, T2 became vulnerable to masking by T3.

Importantly, these "unmasking" studies also speak to the potential influence of attention to the mask on the magnitude of interruption masking. The reduction in attentional availability for T2 caused by the presentation of T1 not only makes T2 vulnerable to masking, but it also reduces the vulnerability of T1 to masking by T2. This cannot be attributed to any change in the availability of attention for T1, because T1 is always the first stimulus to be presented in the sequence, and thus should be fully-attended. Rather the reduction in the magnitude of T1 masking when it is presented with both T2 and T3 must be attributed to the reduction in attentional availability for T2. This suggests that reducing the availability of attention for the mask reduces the magnitude of masking – the opposite of what was found using spatial manipulations of attention. This difference will be discussed in greater detail in the summary section on interruption masking and attention. Interruption Masking and Non-spatiotemporal Manipulations of Attention

# Bottleneck Models of Temporal Processing 87

As can be seen in Table A4, the majority of studies using non-spatiotemporal manipulations of attention have also found that unattended targets are more vulnerable to

-----

Insert Table A4 about here

\_\_\_\_\_

masking. The single exception to this rule comes from the work of Bachmann (1988) in which he found no influence of target-distractor similarity on the magnitude of interruption masking. This finding contrasts with earlier work by Bachmann and Allik (1976), as well as the study of Hines and Smith (1977) which both found highly-similar targets and distractors increased the magnitude of interruption masking. One possible explanation for the difference between these studies is the type of stimuli used in each study. In the case of Bachmann (1988), targets and masks were letters of differing featural similarity. In contrast, in both Bachmann and Allik (1976) and Hines and Smith (1977), stimuli were geometric shapes. It is possible that the use of highly over-learned stimuli such as letters resulted in limited competition for common resources because letters could be processed very efficiently. This would have ameliorated the deficit that would normally have arisen for highly similar targets and masks.

In addition to studies that have used non-spatiotemporal methodologies to vary attention to the target, others have used similar methods to vary attention to the mask. Shelley-Tremblay and Mack (1999; Expt. 3) found much greater impairment when an interruption mask consisted of an observer's own name than when it consisted of a scrambled version of their name, or a high-frequency word. From this, they concluded that the observer's name captured attention away from the target, thereby impairing target processing. This result is similar to that obtained by Hines and Smith (1977). They found that the level of interruption masking depended on whether observers were instructed to attend to the mask. When observers were told to ignore the mask, target identification improved.

Although these findings suggest that attention to the mask does influence the magnitude of interruption masking, the nature of this effect is unclear. In the spatialmanipulation studies reviewed earlier, the results suggested that reducing attentional availability for a mask increases its ability to mask a target by interruption. This is similar to the results obtained using temporal manipulations of attention, but is inconsistent with findings from the non-spatiotemporal studies cited above which indicated that reducing attentional availability for a mask decreases its ability to mask a target by interruption. Interruption Masking and Attention: Summary

The studies that have examined the effects of attentional availability on interruption masking show a clear relationship: unattended targets are more vulnerable to masking. This relationship was found in thirty out of thirty-one studies.

Additionally, there is some evidence that the strength of interruption masking is related to the availability of attention for the mask. This relationship was found in seven out of eight studies. However, the nature of this relationship is unclear. In two studies using spatial manipulations of attention to the mask, increasing the number of masks presented after the target resulted in greater interruption masking. In contrast, in studies using the "unmasking" paradigm, and various non-spatiotemporal manipulations of attention, reducing availability of attention for the mask decreased the magnitude of interruption masking.

There are several potential reasons for this difference. One possibility is that the effects of spatial manipulations of attention on the mask may not have been due to attention, but rather to a change in the level of decision noise or the amount of masking contours. In this case, it would not be surprising if results were different than in the non-spatiotemporal studies where the manipulations are more likely to be influencing only attentional availability. Another possibility is that spatial manipulations of attention have fundamentally different effects than non-spatiotemporal manipulations of attention. This possibility is suggested by the fact that spatial manipulations involve a change in the focus of attention whereas non-spatial manipulations do not. Arguing against this interpretation is the fact that both spatial and non-spatiotemporal manipulations produced similar effects when used to vary attentional availability to the target. At this point, though, before considering any further hypotheses, it would be fruitful to examine the remaining studies in the survey in order to gain further information about the relationship between attention to the mask and the level of masking. The nature of the relationship between attentional availability for the mask and magnitude of masking will be considered again in Section V.

### Integration Masking and Spatial Manipulations of Attention

As can be seen in Table A5, much of the evidence concerning the relationship between integration masking and spatial manipulations of attention comes from studies already cited in

------

Insert Table A5 about here

the section on interruption masking. This is because a number of these studies included conditions with shorter target-mask SOAs that are conducive to integration masking.

With respect to the influence of spatial manipulations of attention on integration masking, the results are clear. None of the studies in the survey indicated that these types of attentional manipulations influence the magnitude of integration masking. This is despite the fact that identical manipulations produced greater levels of interruption masking for unattended targets. This latter point suggests that the null results with respect to integration masking are not simply attributable to an ineffectual manipulation of attention in these experiments.

In contrast to their effect on targets, spatial manipulations of attention to mask displays lead to significant effects on the magnitude of integration masking. In the cases of Cheal and Gregory (1997), Shiu and Pashler (1994) and Tata and Giaschi (1999), greater masking was obtained when the number of masks in the display was increased. The sole exception to this pattern is the findings of Luck, Hillyard, Mouloua, and Hawkins (1996) which indicated no influence of the number of masks on the magnitude of integration masking. As with the results reported on interruption masking, increasing the number of masks in these experiments increased the level of integration masking. This may be interpreted as indicating that unattended masks lead to greater integration masking. However, this interpretation is problematic given other potential explanations such as an increase in decision noise. This issue is discussed at greater length in Section V.

## Integration Masking and Temporal Manipulations of Attention

Studies that have examined this relationship are listed in Table A6. The results suggest that some influence of attention on integration masking may be present. However,

\_\_\_\_\_

Insert Table A6 about here

-----

this influence is relatively small, particularly in comparison to the massive changes in the magnitude of interruption masking that are attributable to attentional availability for a target.

Using an attentional blink paradigm, both Giesbrecht and Di Lollo (1998) and Brehaut et al. (1999) found that the magnitude of integration masking for the second target did not vary as a function of inter-target lag. This suggests that availability of attention for the second target did not influence its vulnerability to integration masking. In a similar experiment, Brehaut et al. (1999) compared the level of integration masking for fullyattended first targets with masking for second targets across a number of inter-target lags. Consistent with the findings of Giesbrecht and Di Lollo (1998), second-target masking was approximately equivalent to the level of first-target masking, and did not vary as a function of lag. This strongly suggests that attention did not influence interruption masking in these experiments.

In contrast to these results, recent work by Visser and Enns (submitted) suggests that integration of stimuli can be influenced by attention. In this study, the attentional blink paradigm was used to vary attentional availability for a dot-matrix integration task. In this task, observers were presented with two frames of dots that when combined yielded a 5x5 matrix of dots. On half the trials, one dot was missing from the completed matrix. In order to accurately report whether a dot was missing, the two frames of dots had to be accurately integrated across an intervening temporal gap – a process that is functionally equivalent to integration masking. Visser and Enns found that at shorter inter-target lags, when attention was less available for the integration task, accuracy in this missing-dot detection task was somewhat reduced. This suggests that attention is required for integration of stimuli, and by extension that integration masking should actually decrease in magnitude when targets are unattended.

## Integration Masking and Non-spatiotemporal Manipulations of Attention

As with the earlier section on integration masking and spatial manipulations of attention, a number of the studies listed in Table A7 were mentioned earlier in the section on

------

Insert Table A7 about here

-----

interruption masking. Again, the rationale for including them in the section on integration masking is because many of them used brief target-mask SOAs.

The results of studies that have used non-spatiotemporal manipulations of attention to the target are unanimous. In each case, unattended targets were more likely to be vulnerable to masking by integration. For example, Shelley-Tremblay and Mack (1999) found a difference in masking level between an observer's own name and a high-frequency word that was maximal at a target-mask SOA of 20 ms. Even more decisive are the results obtained by Bachmann and Allik (1976) and Scheerer (1966). In each of these studies, unattended targets were more vulnerable to masking. Importantly, this vulnerability was greatest when the mask preceded the target. This cannot be attributed to an effect of attention on interruption masking because interruption of target processing cannot occur when the mask comes before the target. Rather, these results must be attributed to an increase in the magnitude of integration masking for unattended targets.

Only one study, by Shelley-Tremblay and Mack (1999), has examined the influence of a non-spatiotemporal manipulation of attention to the masking display. They found that meaningful masks produced greater masking, with a maximal difference between nonmeaningful and meaningful masks occurring at a short target-mask SOA. In contrast to similar studies using spatial manipulations of attention, these results suggest that unattended masks result in a reduction in the magnitude of integration masking.

### Integration Masking and Attention: Summary

The results of studies on the influence of attention on integration masking present a confusing picture. Of the ten studies that have manipulated attentional availability for the target, five have found that integration masking was modulated by attention, and five have not. None of the studies using a spatial manipulation of attention found a relationship to severity of masking. All of the studies using a non-spatiotemporal manipulation did find this relationship. Moreover, although the results from four of the five studies that found a relationship seemed to indicate that unattended targets were more vulnerable to integration masking, the results from one study (Visser & Enns, submitted) seemed to indicate that the opposite was true. Considered collectively, it is impossible to make a strong case either for

۰.

or against a relationship between attention to the target and integration masking. Without additional empirical evidence, the best conclusion that can be made at present is that if such a relationship does exist, the influence of attention on integration masking must be relatively small.

Examination of the studies that have varied attention to the mask suggest that this manipulation does influence the magnitude of integration masking – this is true in four out of five cases. However, the nature of this effect is unclear. In cases in which the number of masks has been varied, ostensibly a spatial manipulation of attention, unattended masks have led to greater masking. In contrast, in the lone example of a non-spatiotemporal manipulation of attention to the mask, unattended masks produced less integration masking. This difference between the effects produced by spatial manipulations and those produced by non-spatiotemporal manipulations echoes those obtained in studies of interruption masking. This point will be discussed further in Section V.

### Lateral Masking and Spatial Manipulations of Attention

Two studies, listed in Table A8, have examined the influence of location cueing on the detection of targets flanked by lateral masks. The results from these studies, however, are

-----

Insert Table A8 about here

------

contradictory. Reinitz (1990) presented a letter target flanked by three distractor letters. This group of four characters could appear either in the top-left, top-right, bottom-right, or bottom-left quadrant of the screen. When the location of the characters was validly cued,

target detection was substantially improved compared to when the location was invalidlycued. This suggests that attending to the target location reduced lateral masking. However, this result differs from the findings of Nazir (1992) who presented targets flanked by masks in one of twelve locations on a notional clockface. In this experiment, identification accuracy did not vary as a function of cue-validity, implying that lateral masking was unaffected by attentional availability for the target.

### Lateral Masking and Temporal Manipulations of Attention

No studies have examined the influence of attention on lateral masking using a temporal manipulation of attention.

### Lateral Masking and Non-spatiotemporal Manipulations of Attention

Several different non-spatiotemporal paradigms have been used in lateral-masking studies. The findings from these different studies, listed in Table A9, have been nearly unanimous

\_\_\_\_\_

Insert Table A9 about here

in indicating that unattended targets are more vulnerable to lateral masking. For example, a number of studies have examined lateral masking as a function of the number of simultaneous flanking masks (e.g. Banks, Larson, & Prinzmetal, 1979; Wolford & Chambers, 1983). It was assumed that when a single target and mask were presented simultaneously, observers would group these two stimuli together, thus reducing attentional availability for the target. In contrast, when the target was presented with multiple masks,

the observer would group the masks together, thus freeing up attentional resources for the target. Consistent with this prediction, when multiple masks flanked a target, lateral masking was much reduced. Equivalently, when more attentional resources were available for target processing, masking was reduced.

The lone exception to this pattern comes from the work of Huckauf, Heller, and Nazir (1999) who varied target-mask similarity in a task in which targets were letters, and masks were either letters or rotated letters. They found no difference in the level of lateral masking between upright and rotated letters, suggesting that attentional availability for the target, which was presumably greater when masks consisted of rotated letters, had no influence on lateral masking. However, it is important to note that like the studies of Bachmann (1986), the experiments of Huckauf et al. (1999) made use of letter targets. Because letters are highly-overlearned stimuli, their identification may have been relatively unimpaired even when unattended. If true, this would likely eliminate any differences between the attended and unattended conditions that might otherwise have occurred with less familiar target stimuli.

## Lateral Masking and Attention: Summary

The findings reviewed above clearly suggest that lateral masking is greater for unattended targets – a result found in six out of eight studies in the survey.

# **SECTION V: General Discussion**

An examination of the studies in this survey clarifies a number of issues about the relationship between attention and masking. However, it predictably raises as many

questions as it answers. What is clear is that availability of attention for the target has a powerful influence on the magnitude of both interruption and lateral masking. In both cases, the vast majority of studies indicate that unattended targets were more vulnerable to both interruption and lateral masking. This evidence comes from a number of different manipulations including visual search, attentional blink, and target-mask similarity.

The results are more equivocal with respect to the relationship between attentional availability for the target and the level of integration masking. The studies in the survey were split down the middle, with half showing a relationship, and half showing no relationship. Moreover, among the studies that did find a relationship, while the majority found that unattended targets were more vulnerable to integration masking, one study found the opposite to be true. Nonetheless, it seems possible that some small effects of attention on integration masking may be present. At the very least, further empirical investigation is warranted.

Of particular interest would be an investigation of possible differences between spatial and non-spatiotemporal manipulations of attention as they relate to integration masking. The literature reviewed above exhibits a clear distinction between these two categories with spatial manipulations consistently failing to influence integration masking, and non-spatiotemporal manipulations always influencing integration masking. This suggests that the vulnerability of targets to integration may depend on the availability of certain types of attentional resources. For example, it may be that whether a target is at the focus of the attentional "spotlight" does not influence integration, whereas the ability of a target to capture processing resources and thus be processed quickly does influence integration. This suggestion is consistent with a number of studies which indicate that spatial location plays a unique role in visual processing (e.g. Visser, Bischof & Di Lollo, 1998). <u>The effect of attentional availability for the mask</u>

The review of the literature suggests that attentional availability for processing of the mask may also play a role in the magnitude of different types of masking. Indeed, summed across the different types of attentional manipulations and mask types, a total of eleven out of thirteen studies have found evidence consistent with this proposition. Broken down by type of pattern masking, the evidence is strongly in favour of an effect of attention to the mask on both the level of interruption and integration masking. As for lateral masking, given that the targets and masks appear simultaneously in this paradigm, it is not possible to meaningfully separate attentional manipulations to the mask from those to the target. For example, increasing the number of flanking masks in the experiments of Banks et al. (1979) can be interpreted as both increasing the availability of attention for the target, via grouping, and as decreasing the availability of attention for the masks via set-size manipulation. Thus, the following discussion will consider only interruption and integration masking.

As noted in Section IV, while there is an effect of attention to the mask on the magnitude of masking, the exact nature of this effect is unclear. On one hand, increasing the number of masks appears to increase the magnitude of both interruption and integration masking. Interpreted from within the traditional visual search framework, this finding suggests that when attention is divided by increasing the number of masks, the vulnerability of targets to masking is increased. However, these results are the opposite of those found using other temporal and non-spatiotemporal paradigms. Here, manipulations of attention

such as requiring observers to either report the identity of the mask, or ignore the mask (Hines & Smith, 1977), indicated that when attentional availability for the mask is decreased, both interruption and integration masking are correspondingly ameliorated.

One solution to this inconsistency may lie in a reinterpretation of the evidence from studies that have varied the number of masks following a target. As noted in Section IV, the effect of this increase may not be limited to a division of attention. It may also produce an increase in decision noise by creating confusion as to the correct location of a target. This increase in decision noise, which is not present in temporal or non-spatiotemporal manipulations of attention, may have overpowered any reduction in vulnerability to masking that resulted from a division of attention to the masks. One way to test this proposition is by presenting displays that consist of both targets and distractors, and either a single mask at the target location, or multiple masks at all locations. Using this paradigm, an equivalent amount of decision noise would be present on each trial due to the presence of the distractors. Therefore, any influence of the number of masks on the strength of masking could be attributed unambiguously to attentional factors.

Another possibility is that introducing multiple masks simply increases the amount of masking contour, thus increasing the amount of masking. Once again, this increase in the amount of contours may produce additional target masking that overwhelms any benefits of dividing attention to the mask. One way to test this hypothesis is by comparing displays that are masked by one small mask, four small masks, or one large mask with an equivalent amount of contours to the four small masks. If the increase in masking produced by increasing the number of masks is due to the amount of masking contours than equivalent

masking should be found when using either one large mask or four small masks. In contrast if the increase in masking is attributable to a division of attention between the masks, only the condition with four masks should produce this increase because attention is not divided when one large mask is used.

This analysis clearly suggests that an interpretation of the results from studies that have varied the number of masks will not be possible without further experiments to clarify the nature of the effects produced by this manipulation. However, the remaining evidence from temporal and non-spatiotemporal manipulations of attention suggests that there is clearly an influence of attentional availability for the mask that remains to be accounted for in any comprehensive theory of masking.

In the next section, I begin by outlining aspects of such a theory based on the object substitution account of masking proposed by Di Lollo et al. (in press; see also Enns & Di Lollo, 1997; Giesbrecht & Di Lollo, 1998). This model was created to explain the influence of attention on interruption masking. Thus, the parameters were not designed to account for attentional effects on lateral masking, or for the influence of attention the mask on interruption and integration masking. Nonetheless, with very few modifications, this framework may be expanded to account for these findings. I begin by outlining the basic framework of the object-substitution model. Then, in subsequent sections, I outline modifications to the model to allow it to account for attentional effects on lateral masking, as well as the influence of attended and unattended masks.

A theory of masking by object substitution

The majority of theories of visual perception view the process as feed-forward. That is, visual perception is seen as progressing through a series of hierarchical stages. This begins with transduction of light into electrical impulses in the eye, continues through to the decomposition of the image into simple features in V1, and then ends with a re-assembly of these simple features into objects in higher brain regions. Importantly, progress through the stages is viewed as unidirectional, with information being passed from lower stages to higher stages, but not vice-versa.

With respect to theories of masking, this viewpoint has led to several vexing questions. Chief amongst these is the issue of how neural impulses corresponding to the mask can "catch up" to neural impulses of the target when the target is presented prior to the mask. Viewed from a feed-forward perspective, such "catch-up" would be the only possible means for a trailing mask to interfere with target processing. To answer this question, Breitmeyer and Ganz (1976; Breitmeyer, 1984) suggested a model based on transient and sustained pathways in vision. These two pathways are known to be anatomically and functionally separate. The transient pathway carries information about rapid onsets of stimuli, and subserves motion perception mechanisms. The sustained pathway carries information about stimuli that remain visible for a period of time, and subserves object recognition. More importantly, information is known to travel faster along the transient pathway than the sustained pathway. On the basis of these characteristics, Breitmeyer and Ganz (1976) argued that backward masking arises when activation in the sustained pathway, carrying information necessary for target identification, is impaired by activation carried along the transient pathway, signalling the onset of the mask. Such interference is possible
because transient signals are transmitted faster than sustained signals, thus allowing a neural representation of the mask to "catch up" to the representation of the prior target.

Although this conceptualization of backward masking does solve the problem of how a mask presented after a target can impair its identification, the theory cannot account for more recent findings. In their four-dot masking paradigm, Enns & Di Lollo (1997) presented a target and a mask simultaneously – the target consisted of a Landolt C while the mask consisted of four surrounding dots. When the target and mask offset simultaneously, target identification was unimpaired. However, when the mask remained on the screen after target offset, impairment began to occur. Moreover, the longer the duration that the mask remained on the screen, the greater the level of impairment.

This finding cannot be accounted for by transient-on-sustained inhibition because, being simultaneously presented with the target, the mask does not generate a separate transient signal. Moreover, the fact that the magnitude of masking increases with the duration of the mask also argues against some sort of inhibition resulting from the offset of the target. In truth, the problem of explaining four-dot masking is troublesome for any feedforward theory of masking because interference is not produced by the onset of a "new" stimulus. Rather, parts of a single stimulus begin to interfere with each other when some parts remain in view after others have disappeared. It is difficult to see how any feedforward theory could predict such a phenomenon without invoking a complex series of operations.

Because of the difficulties of creating a feed-forward model of the four-dot masking results, Di Lollo et al. (in press) proposed a theory of masking based on a fundamentally

different view of visual processing. Rather than a hierarchical system of processes that feeds information only in one direction, Di Lollo et al. (in press) built their theory of <u>masking by</u> <u>object substitution</u> on the principle of re-entrant processing in which information flows in both directions between higher and lower levels of the visual system. Evidence for the existence of re-entrant pathways in the visual system is plentiful (Felleman & Van Essen, 1991; Perkel, Bullier, & Kennedy, 1986; Shipp & Zeki, 1989; Sillito, Jones, Gerstein, & West, 1994; Zeki, 1993). However, only recently have theories began to make use of these pathways in explanations of empirical phenomena (e.g. Grossberg, 1995; Mumford, 1992).

To explain backward masking in general, and four-dot masking in particular, Di Lollo et al. (in press) suggested that processing of a target stimulus begins with an initial decomposition into simple features at low levels in the visual system (e.g. V1). This information then cascades forward to higher areas in the visual system which generate <u>perceptual hypotheses</u> about the identity of the target. For example, given a target letter "T", potential perceptual hypotheses about the target's identity may include not only "T' but also stimuli with similar features such as "L" or "7". In order for target identification to occur, a perceptual hypothesis must be correlated with available information about a stimulus in lower visual areas – a process referred to as <u>iterative correlation</u>. This iterative correlation process continues for a number of cycles, with information being continuously fed from lower to higher visual areas and back again, until a critical correlation level is achieved. Having achieved this critical correlation, the perceptual hypothesis is confirmed and the object is identified.

Within this framework, masking occurs when a representation of the mask overwrites the representation of the target in lower visual areas before the perceptual hypothesis about the target's identity has been confirmed. As a result of this overwriting, a new perceptual hypothesis is created that corresponds to the identity of the mask. This results in perception of the mask alone, with the only evidence that the target was ever presented being limited to instances of subliminal priming (see Marcel, 1983; Visser, Merikle, & Di Lollo, 1998). This theory provides a ready account of the existing literature on backward masking. It can also account for four-dot masking if it is assumed that the perceptual hypothesis about the target surrounded by the four dots cannot be confirmed before the target is removed from the display. Under these circumstances, the target representation begins to rapidly decay in V1, while the representation of the four dots remaining on the screen is unaffected. As a result, the support for the perceptual hypothesis corresponding to the target decreases, while support for the perceptual hypothesis corresponding to only the four dots begins to increase. Ultimately, the longer the dots remain on the screen by themselves, the more likely that a perceptual hypothesis corresponding to the dots alone will be confirmed. This nicely accounts for the increase in masking found when the four dots are allowed to remain on the screen for a longer duration.

For the present purposes, an important question is how the object-substitution theory of masking can explain the influence of attention on masking. Di Lollo et al. (in press) suggested that attention modulates the speed of the iterative correlation process. When attention can be focussed on a single target, such as when a visual search display contains no distractors, verification of a perceptual hypothesis can be carried out quickly. This increases the likelihood that a target will be identified before the onset of the trailing mask can disrupt the correlation process. In contrast, when attention cannot be focussed on a target, the iterative correlation process is slowed. This increases the likelihood that the mask will be presented before the target's perceptual hypothesis can be verified.

This framework provides a parsimonious account of the interruption masking results reviewed in the present work. As noted in Section IV, unattended targets were found to be more vulnerable to interruption masking than attended targets. Viewed from the perspective of object-substitution, this suggests that for unattended targets, verification of target identity was slowed, resulting in greater interference from the temporally-trailing interruption mask. In contrast, for attended targets, verification of target identity was relatively faster, allowing it to escape interference from the trailing interruption mask.

## Applying object substitution to other forms of masking

Although the object-substitution framework of Di Lollo et al. (in press) accounts well for situations in which the mask remains visible after the target has disappeared (i.e. backward masking and four-dot masking), additional considerations are necessary to account for other types of masking such as lateral and integration masking. This is because these types of masking are most effective when the mask appears simultaneously with the target. These conditions obviously preclude an explanation that relies on mask activation exceeding target activation in lower visual areas.

Nonetheless, there are a number of possible ways that integration and lateral masking can be accounted for within the existing object-substitution framework. The most likely

scenario is that integration and lateral masking interfere with the initial representation of the target at low-levels of the visual system. With sufficient degradation of the low-level target representation, verification of a perceptual hypothesis about its identity may become impossible. This explanation is consistent with existing accounts of integration and lateral masking (see Turvey, 1973; Wolford & Chambers, 1983). It is also consistent with phenomenological evidence (Liss, 1968) showing that targets masked by simultaneous stimuli are present long enough to be identified, but that identification is simply made impossible by the presence of the masking contours.

Given this hypothetical locus of integration and lateral masking effects within the object-substitution framework, what remains to be considered is how attentional availability for a target could modulate its vulnerability to lateral and integration masking. An account of these attentional effects is given in the following section. This account focuses on lateral masking because of the strong evidence that it is influenced by attention. However, some mention is also made of integration masking because there is at least some suggestion that attentional availability may play a small role in integration masking.

## Lateral masking and Attention

As noted in Section II, several explanations for lateral masking have been advanced, ranging from response interference (e.g. Eriksen, O'Hara, & Eriksen, 1982) to contour interaction (Bouma, 1970). Given the nature of the stimuli used in the studies listed in Section IV, which were generally unlikely to lead to response interference, the effects of attention on lateral masking seem to be primarily attributable to low-level contour

interactions. For this reason, the explanation advanced here is couched entirely in terms of contour interference.

One factor likely to influence lateral masking is the spatial resolution at which a target is encoded. The higher the spatial resolution, the less likely that contours from an adjacent mask would be confused with the target. This suggests that the influence of attentional availability for a target on lateral masking may lie in the influence of attention on spatial resolution. Specifically, unattended targets may be coded with lower spatial resolution, thus making them more vulnerable to interference from adjacent mask contours. This possibility was also considered by Enns and Di Lollo (1997) who suggested that four-dot masking may also be partially due to a reduction in spatiotemporal resolution of unattended stimuli.

More direct evidence in support of the link between attention and spatial resolution comes from the work of Yeshurun and Carrasco (1999) who examined vernier acuity as a function of location cueing. They found that when the location of a vernier target was validly cued that accuracy increased and response times were faster than when the location was invalidly cued. Based on these results, the authors argued that focussing attention on the location of a target stimulus increased the spatial resolution of the target representation, thereby improving task performance.

Couched in terms of the object substitution framework, an increase in the spatial resolution of the target at low levels in the visual system would be likely to increase the probability that the correct perceptual hypothesis about the target's identity would be verified. That is, by increasing the quality of a target's representation at a low-level in the

visual system, the perceptual hypothesis made about the target's identity is more likely to be accurate. This can be likened to the process of recognizing handwriting. If we are attempting to read a word written in a barely legible script, perhaps like that of a medical doctor, we will likely have to take several guesses before correctly identifying a word. Moreover, if the quality is poor enough, we may never identify the word. In contrast, if we are attempting to read very precise, tidy script, it is likely that we will be able to identify the words correctly on the first attempt, with almost no instances in which a word will be completely unidentifiable.

It should be noted that an analogous account can be made for effects of attentional availability on the vulnerability of a target to integration masking. A number of lines of evidence suggest that attended objects may be coded with greater stimulus quality. For example, it is well known that attended objects appear to be phenomenologically clearer and sharper than unattended ones (James, 1890/1950). Moreover, attended objects are perceived to last for a longer duration (Enns, Brehaut, & Shore, 1999; Mattes & Ulrich, 1998). As with the case of lateral masking, an increase in the stimulus quality of a target may make it less vulnerable to interference with overlapping contours from an integration mask. As a result, the probability of confirming a perceptual hypothesis that corresponded to the target's identity would be increased.

# The Influence of Attention to the Mask

The results from a number of studies listed in Section IV suggest that both integration and interruption masking can be significantly influenced by allocation of attention to the mask. Setting aside the results from studies that have manipulated the number of masks, the remaining studies suggest that attended masks produce greater interruption and integration masking, while such masking is reduced when masks are unattended. One possible explanation for the effects on interruption masking emerges from a closer consideration of the process of target identification that occurs in a typical masking experiment.

Given a display in which a target and mask are presented in rapid sequence, the task for the observer is to decide the identity of the stimulus presented at the target location. Under these conditions, there are two possibilities – the target or the mask. Similarly, in the context of the object substitution model, there are two possible perceptual hypotheses to be confirmed – a hypothesis corresponding to the target, or a hypothesis corresponding to the mask. The process underlying identification of a masked target can thus be reduced to a competition between two possible perceptual hypotheses.

In such a competition, the stimulus that requires the fewest iterations to have its perceptual hypothesis confirmed has an advantage. With respect to the influence of attention in this competition, Di Lollo and colleagues have already hypothesized that the confirmation of perceptual hypotheses for unattended targets is slower. Thus, it makes sense to assume that the process of confirming the perceptual hypothesis about the identity of an unattended mask would also be slower. Given this assumption, it follows that in a competition between an attended target and an unattended mask, some advantage would be conferred to the target by the additional time needed to verify the perceptual hypothesis corresponding to the mask's identity.

This explanation is particularly suited to the results of Shelley-Tremblay and Mack (1999) which showed that a mask composed of an observer's own name was more effective

than a scrambled version of their name. Because, a perceptual hypothesis corresponding to an observer's name is likely to require few iterations to be verified, it is likely to win a competition with a hypothesis corresponding to a less-meaningful target stimulus. As a result, a meaningful mask will produce greater interruption masking than a less meaningful one.

A similar explanation can be advanced to explain the effects of attention to the mask on integration masking. An additional effect may occur at a lower level in the visual system. Just as attended targets are likely to be coded with greater stimulus quality, attended masks will receive similar benefits. Thus in the case where target and mask are integrated at a lowlevel in the visual system, attended masks will produce more masking than unattended masks because their higher stimulus quality will lead to greater degradation of the target representation. This makes successful confirmation of a perceptual hypothesis corresponding to the target's identity less likely.

#### Concluding Comments

In the present work, an analysis of studies on attention and masking has revealed that attention to masked targets modulates their vulnerability to interruption and lateral masking. In addition, attention to the mask also appears to modulate the level of interruption and integration masking. Beyond these conclusions, several issues have emerged that are in need of further empirical investigation.

One area is the influence of attention on integration masking. Existing studies on this question have primarily concerned interruption masking, with information on integration masking gleaned from a few conditions in which short target-mask SOAs were employed.

What is needed are studies that focus exclusively on the influence of attention on integration masking. A second area warranting investigation is the influence of spatial manipulations of attention to masks. To date, none of the studies that have touched upon this issue have used manipulations that have clearly varied only attentional availability to masking displays.

Although much empirical work remains to be done, the conclusions of this study clearly suggest that attention and masking play important roles in visual processing. Attended targets are generally invulnerable to a number of different types of pattern masks. In contrast, unattended targets are vulnerable to a number of different types of masking. This suggests that masking functions at a number of different levels in the stream of visual processing to suppress unattended stimuli. This is consistent with the role of masking as a mechanism that replaces less biologically relevant stimuli with new visual input, thereby increasing the efficiency of visual processing.

## Appendix A: References

Alpern, M. (1953). Metacontrast. Journal of the Optical Society of America, 43, 648-657.

- Arand, D., & Dember, W. N. (1976). Metacontrast with internal contours in target and mask. Bulletin of the Psychonomic Society, 7, 370-372.
- Averbach, E., & Sperling, G. (1961). Short-term storage of information in vision. In C.Cherry (Ed.), <u>Symposium on information theory</u> (pp. 196-211). London:Butterworth.
- Bachmann, T. (1988). Time course of the subjective contrast enhancement for a second stimulus in successively paired above-threshold transient forms: Perceptual retouch instead of forward masking. <u>Vision Research</u>, 28, 1255-1261.
- Bachmann, T., & Allik, J. (1976). Integration and interruption in the masking of form by form. <u>Perception</u>, 5, 79-97.
- Banks, W. P., Larson, D. W., & Prinzmetal, W. (1979). Asymmetry of visual interference. Perception and Psychophysics, 25, 447-456.
- Banks, W. P., & White, H. (1984). Lateral interference and perceptual grouping in visual detection. <u>Perception and Psychophysics</u>, <u>36</u>, 285-295.
- Beck, J. & Ambler, B. (1972). Discriminability of differences in line slope and in line arrangement as a function of mask delay. <u>Perception and Psychophysics</u>, 12, 33-38.

Bouma, H. (1970). Interaction effects in parafoveal letter recognition. Nature, 226, 177-178.

Braddick, O. (1973). The masking of apparent motion in random-dot patterns. <u>Vision</u> <u>Research</u>, <u>13</u>, 355-369.

- Brehaut, J., Enns, J. T., & Di Lollo, V. (1999). Visual masking plays two roles in the attentional blink. <u>Perception and Psychophysics</u>, <u>61</u>, 1436-1448.
- Breitmeyer, B. G. (1984). <u>Visual masking: An integrative approach.</u> New York: Oxford University Press.
- Breitmeyer, B.G., & Ganz, L. (1976). Implications of sustained and transient channels for theories of visual pattern masking, saccadic suppression, and information processing.
   <u>Psychological Review</u>, <u>83</u>, 1-36.
- Broadbent, D. E, & Broadbent, M. H. (1987). From detection to identification: Response to multiple targets in rapid serial visual presentation. <u>Perception and Psychophysics</u>, <u>42</u>, 105-113.
- Cheal, M., & Gregory, M. (1997). Evidence of limited capacity and noise reduction with single-element displays in the location-cuing paradigm. Journal of Experimental <u>Psychology: Human Perception and Performance</u>, 23, 51-71.
- Chun, M. M., & Potter, M. C. (1995). A two-stage model for multiple target detection in rapid serial visual presentation. <u>Journal of Experimental Psychology: Human</u> <u>Perception and Performance</u>, <u>21</u>, 109-127.
- Cohene, L. S., & Bechtoldt, H. P. (1974). Visual recognition as a function of stimulus offset asynchrony and duration. <u>Perception and Psychophysics</u>, <u>15</u>, 221-226.
- Cohene, L. S., & Bechtoldt, H. P. (1975). Visual recognition of dot-pattern bigrams: An extension and replication. <u>American Journal of Psychology</u>, <u>88</u>, 187-199.

- Debner, J. A., & Jacoby, L. L. (1994). Unconscious perception: Attention, awareness, and control. <u>Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,</u> 304-317.
- Dember, W. N., & Purcell, D. G. (1967). Recovery of masked visual targets by inhibition of the masking stimulus. <u>Science</u>, <u>157</u>, 1335-1336.
- Di Lollo, V., Enns, J. T., & Rensink, R. (in press). Psychophysics of re-entrant pathways: A computational model of object substitution (CMOS). Journal of Experimental <u>Psychology: General</u>.
- Enns, J. T., Brehaut, J. C., & Shore, D. I. (1999). The duration of a brief event in the mind's eye. <u>The Journal of General Psychology</u>, <u>126</u>, 355-372.
- Enns, J. T., & Di Lollo, V. (1997). Object substitution: A new form of masking in unattended visual locations. <u>Psychological Science</u>, <u>8</u>, 135-139.
- Eriksen, C. W. (1966). Temporal luminance summation effects in backward and forward masking. <u>Perception and Psychophysics</u>, <u>1</u>, 87-92.
- Eriksen, C. W., & Rohrbaugh, J. (1970). Visual masking in multielement displays. Journal of Experimental Psychology, 83, 147-154.
- Eriksen, C. W., & Collins, J. F. (1969). Visual perceptual rate under two conditions of search. Journal of Experimental Psychology, 80, 498-492.
- Eriksen, C. W., O'Hara, W. P., & Eriksen, B. (1982). Response competition effects in samedifferent judgments. <u>Perception and Psychophysics</u>, 32, 261-270.
- Eriksen, C. W., & Yeh, Y. (1985). Allocation of attention in the visual field. Journal of Experimental Psychology: Human Perception and Performance, <u>11</u>, 583-597.

- Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in primate visual cortex. <u>Cerebral Cortex</u>, <u>1</u>, 1-47.
- Fisher, D. L., Duffy, S. A., Young, C., & Pollatsek, A. (1988). Understanding the central processing limit in consistent-mapping visual search tasks. <u>Journal of Experimental</u> <u>Psychology: Human Perception and Performance, 14</u>, 253-266.
- Giesbrecht, B., & Di Lollo, V. (1998). Beyond the attentional blink: Visual masking by object substitution. Journal of Experimental Psychology: Human Perception and Performance, 24, 1454-1466.
- Grossberg, S. (1995). 3-D vision and figure-ground separation by visual cortex. <u>Perception</u> <u>and Psychophysics</u>, <u>55</u>, 48-120.
- Harcum, E. R., & Shaw, M. R. (1974). Cognitive and sensory lateral masking of tachistoscopic patterns. Journal of Experimental Psychology, 103, 663-667.
- Henderson, J. M. (1991). Stimulus discrimination following covert attentional orienting to an exogenous cue. Journal of Experimental Psychology: Human Perception and Performance, 17, 91-106.
- Hines, D. & Smith, S. (1977). Recognition of random shapes followed at varying delays by attended or unattended shapes, digits, and line grids. <u>Journal of Experimental</u>
  <u>Psychology: Human Learning and Memory, 3</u>, 29-36.
- Huckauf, A., Heller, D., & Nazir, T. A. (1999). Lateral masking: Limitations of the feature interaction account. <u>Perception and Psychophysics</u>, <u>61</u>, 177-189.
- Jacobson, J. Z. (1974). Interaction of similarity to words of visual masks and targets. Journal of Experimental Psychology, 102, 431-434.

- Jacoby, L. L., Toth, J. P., & Yonelinas, A. P. (1993). Separating conscious and unconscious influences of memory: Measuring recollection. <u>Journal of Experimental Psychology:</u> <u>General, 122, 139-154</u>.
- James, W. (1950). <u>The principles of psychology Vol. I.</u> New York: Dover (original published 1890).
- Jolicoeur, P. (1998). Modulation of the attentional blink by on-line response selection: Evidence from speeded and unspeeded Task-1 decisions. <u>Memory and Cognition</u>, <u>26</u>, 1014-1032.
- Kahneman, D. (1967). An onset-onset law for one case of apparent motion and metacontrast. <u>Perception and Psychophysics</u>, 2, 557-583.
- Kolers, P. A. (1968). Some psychological aspects of pattern recognition. In P. A. Kolers andM. Eden (Eds.) <u>Recognizing Patterns</u>. (pp. 4-61). Boston: MIT Press.
- Krose, B. J. A., & Julesz, B. (1989). The control and speed of shifts of attention. <u>Vision</u> <u>Research</u>, <u>29</u>, 1607-1619.
- Liss, P. (1968). Does backward masking by visual noise stop stimulus processing? <u>Perception and Psychophysics</u>, <u>4</u>, 328-330.
- Luck, S. J., Hillyard, S. A., Mouloua, M., & Hawkins, H. L. (1996). Mechanisms of visual-spatial attention: Resource allocation or uncertainty reduction? <u>Journal of</u> <u>Experimental Psychology: Human Perception and Performance</u>, <u>22</u>, 725-737.

Mack, A. & Rock, I. (1998). Inattentional blindness. Cambridge: MIT Press.

Marcel, A. J. (1983). Conscious and unconscious perception: Experiments on visual masking and word recognition. <u>Cognitive Psychology</u>, 15, 197-237.

- Mattes, S., & Ulrich, R. (1998). Directed attention prolongs the perceived duration of a brief stimulus. <u>Perception and Psychophysics</u>, <u>60</u>, 1305-1317.
- Merikle, P. M. (1977). On the nature of metacontrast with complex targets and masks. Journal of Experimental Psychology: Human Perception and Performance, <u>3</u>, 607-621.
- Michaels, C. F., & Turvey, M. T. (1979). Central sources of visual masking: Indexing structures supporting see at a single, brief glance. <u>Psychological Research</u>, <u>41</u>, 1-61.
- Mumford, D. (1992). On the computational architecture of the neocortex II: The role of cortico-cortical loops. <u>Biological Cybernetics</u>, <u>66</u>, 241-251.
- Nazir, T. A. (1992). Effects of lateral masking and spatial precueing on gap-resolution in central and peripheral vision. <u>Vision Research</u>, <u>32</u>, 771-777.
- Perkel, D. J., Bullier, J., & Kennedy, H. (1986). Topography of the afferent connectivity of area 17 in the Macaque monkey. Journal of Comparative Neurology, 253, 347-402.
- Posner, M. I. (1980). Orienting of attention. <u>Quarterly Journal of Experimental Psychology</u>, <u>32</u>, 3-25.
- Nakayama, K., & Mackeben, M. (1989). Sustained and transient components of focal visual attention. <u>Vision Research</u>, 29, 1631-1647.
- Ramachandran, V. S., & Cobb, S. (1995). Visual attention modulates metacontrast masking. Nature, <u>373</u>, 66-68.
- Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992). Temporary suppression of visual processing of visual processing in an RSVP task: An attentional blink? <u>Journal of</u> <u>Experimental Psychology: Human Perception and Performance</u>, <u>18</u>, 849-860.

- Reinitz, M. T. (1990). Effects of spatially directed attention on visual encoding. <u>Perception</u> <u>and Psychophysics</u>, 47, 497-505.
- Schiller, P. H. (1965). Monoptic and dichoptic visual masking by patterns and flashes. Journal of Experimental Psychology, 69, 193-199.
- Schiller, P. H. (1966). Forward and backward masking as a function of relative overlap and intensity of test and masking stimuli. <u>Perception and Psychophysics</u>, <u>1</u>, 161-164.
- Shapiro, K. L., Raymond, J. E., & Arnell, K. M. (1994). Attention to visual pattern information produces the attentional blink in RSVP. <u>Journal of Experimental</u> <u>Psychology: Human Perception and Performance, 20</u>, 357-371.
- Shelley-Tremblay, J. & Mack, A. (1999). Metacontrast masking and attention. <u>Psychological</u> <u>Science</u>, <u>10</u>, 508-515.
- Shipp, S., & Zeki, S. (1989). The organization of connections between areas V5 and V1 in the macaque monkey visual cortex. <u>European Journal of Neuroscience</u>, <u>1</u>, 309-332.
- Shiu, L., & Pashler, H. (1994). Negligible effect of spatial precuing on identification of single digits. <u>Journal of Experimental Psychology: Human Perception and</u> <u>Performance, 20</u>, 1037-1054.
- Sillito, A. M., Jones, H. E., Gerstein, G. L., & West, D. C. (1994). Feature-linked synchronization of thalamic relay cell firing induced by feedback from visual cortex. <u>Nature, 369</u>, 479-482.
- Smith, M. C., & Schiller, P. H. (1966). Forward and backward masking: a comparison. Canadian Journal of Psychology, 20, 337-342.

- Spencer, T. J. (1969). Some effects of different masking stimuli on iconic storage. Journal of Experimental Psychology, 81, 132-140.
- Spencer, T. J. (1972). Backward masking: Stimulus energy and positional uncertainty. American Journal of Psychology, 85, 535-547.
- Spencer, T. J., Hawkes, L., & Mattson, G. (1972). Effect of a forward indicator on backward masking. Journal of Experimental Psychology, 93, 297-301.
- Spencer, T. J., & Shuntich, R. (1970). Evidence for an interruption theory of backward masking. <u>Journal of Experimental Psychology</u>, <u>85</u>, 198-203.
- Stoper, A. E., & Banffy, S. (1977). Relation of split apparent motion to metacontrast. Journal of Experimental Psychology: Human Perception and Performance, <u>3</u>, 258-277.
- Tata, M. S., Giaschi, D. E., & Di Lollo, V. (1998). Visual attention modulates metacontrast masking. Poster presented at the annual meeting of the Association for Research in Vision and Ophthalmology, Ft. Lauderdale, FL.
- Tata, M. S., & Giaschi, D. E. (1999). Selective attention to the mask modulates object substitution masking. Poster presented at the annual meeting of the Association for Research in Vision and Ophthalmology, Ft. Lauderdale, FL.
- Turvey, M. T. (1973). On peripheral and central processes in vision: inferences from an information-processing analysis of masking with patterned stimuli. <u>Psychological</u> <u>Review</u>, <u>80</u>, 1-52.

- Visser, T. A. W., Bischof, W. F., & Di Lollo, V. (1999). Attentional switching in spatial and non-spatial domains: Evidence from the attentional blink. <u>Psychological Bulletin</u>, <u>125</u>, 458-469.
- Visser, T. A. W., & Enns, J. T. (submitted). The role of attention in temporal integration. <u>Perception</u>.
- Visser, T. A. W., Merikle, P. M., & Di Lollo, V. (1998, November). A tale of two masks: Linking unconscious perception, the attentional blink, and masked priming. Paper presented at the 39<sup>th</sup> Annual Meeting of the Psychonomic Society, Dallas, TX.
- Visser, T. A. W., Zuvic, S. M., Bischof, W. F., & Di Lollo, V. (1999). The attentional blink with targets in different spatial locations. <u>Psychonomic Bulletin and Review, 6</u>, 432-436.
- Weisstein, N. (1966). Backward masking and models of perceptual processing. Journal of Experimental Psychology, 72, 232-240.
- Weisstein, N. (1972). Metacontrast. In D. Jameson & L. M. Hurvich (Eds.) <u>Handbook of sensory physiology, Vol. 7/4, Visual psychophysics</u> (pp. 233-272). New York: Springer-Verlag.
- Wolford, G., & Chambers, L. (1983). Lateral masking as a function of spacing. <u>Perception</u> and <u>Psychophysics</u>, <u>33</u>, 129-138.
- Yeshurun, Y., & Carrasco, M. (1999). Spatial attention improves performance in spatial resolution tasks. <u>Vision Research</u>, <u>39</u>, 293-306.
- Zeki, S. (1993). A vision of the brain. Oxford: Blackwell.

		Type of Attentional Manipulation			
		Non-spatio temporal	Spatial	Temporal	
Type of Masking	Integration	4	7	3	
	Interruption	9	17	5	
	Lateral	6	2	0	

**Table A1.** Number of published studies as a function of type of masking and attentional manipulation. The grayed-out box indicates that no published studies have examined the influence of a manipulation of temporal attention on lateral masking.

Author(s)	Type of Attentional Manipulation		Influence on Masking	
	Target	Mask	Attention to Target	Attention to Mask
Beck & Ambler (1972)	Visual Search		Yes	
Cohene & Bechtoldt (1974)	Visual Search		Yes	
Cohene & Bechtoldt (1975)	Visual Search		Yes	
Di Lollo, Enns, & Rensink	Visual Search /		Yes	
(in press)	Location Cueing			
Enns & Di Lollo (1997)	Visual Search		Yes	
Eriksen & Rohrbaugh (1970)	Visual Search		Yes	
Fisher, Duffy, Young, &	Visual Search		Yes	
Pollatsek (1988)				
Henderson (1991)	Location Cueing		Yes	
Krose & Julesz (1989)	Location Cueing		Yes	
Shiu & Pashler (1994)	Location Cueing	Visual Search	Yes	Yes
Spencer (1969)	Visual Search		Yes	
Spencer (1972)	Visual Search		Yes	
Spencer, Hawkes, & Mattson	Location Cueing		Yes	
(1972)				
Spencer & Shuntich (1970)	Visual Search		Yes	
Tata, Giaschi, & Di Lollo	Visual Search /		Yes	
(1998)	Location Cueing			
Tata & Giaschi (1999)	Visual Search	Visual Search	Yes	Yes
Weisstein (1966)	Visual Search		Yes	

**Table A2.** Influence of spatial manipulations of attention on interruption masking. The text in the second and third columns indicate whether attention was varied to the target or to the mask, and the type of attentional manipulation: "Visual search" indicates a target or mask was embedded amongst a variable number of distractors; "Location Cueing" indicates that the location of a target or mask embedded amongst distractors was indicated before onset of a search display; "---" indicates no manipulation of that type was made in the experiment. The text in the fourth and fifth columns indicates whether a spatial manipulation of attention influenced the level of interruption masking: "Yes" indicates that there was an effect of attention; "No" indicates that there was no effect of attention; "----" indicates that no manipulation of that type was made in the experiment.

Author(s)	Type of Attentional Manipulation		Influence on Masking	
	Target	Mask	Attention to Target	Attention to Mask
Brehaut, Enns, & Di Lollo (1999)	Attentional blink		Yes	
Dember & Purcell (1967)	Unmasking	Unmasking	Yes	Yes
Giesbrecht & Di Lollo (1998)	Attentional blink		Yes	
Michaels & Turvey (1979)	Unmasking	Unmasking	Yes	Yes
Turvey (1973)	Unmasking	Unmasking	Yes	Yes

**Table A3.** Influence of temporal manipulations of attention on interruption masking. The text in the second and third columns indicate whether attention was varied to the target or to the mask, and the type of attentional manipulation: "Attentional blink" indicates the attentional blink paradigm (described in text) was used to vary attention; "unmasking" indicates an analagous paradigm that uses three target stimuli presented in rapid succession (described in text); "---" indicates no manipulation of that type was made in the experiment. The text in the fourth and fifth columns indicates whether a temporal manipulation of attention influenced the level of interruption masking: "Yes" indicates that there was an effect of attention; "No" indicates that there was no effect of attention; "---" indicates that no manipulation of that type was made in the experiment.

Author(s)	Type of Attentional Manipulation		Influence on Masking	
	Target	Mask	Attention to Target	Attention to Mask
Arand & Dember (1976)	Target-mask similarity		Yes	
Bachmann (1988)	Target-mask similarity		No	
Bachmann & Allik (1976)	Target-mask similarity		Yes	
Eriksen & Collins (1969)	Attentional set		Yes	
Hines & Smith (1977)	Target-mask similarity	Attentional set	Yes	Yes
Merikle (1977)	Target-mask similarity		Yes	
Michaels & Turvey (1979)	Meaningfulness	Meaningfulness	Yes	No
Ramachandran & Cobb (1995)	Grouping		Yes	
Shelley-Tremblay & Mack (1999)	Meaningfulness	Meaningfulness	Yes	Yes

**Table A4.** Influence of non-spatiotemporal manipulations of attention on interruption masking. The text in the second and third columns indicate whether attention was varied to the target or to the mask, and the type of attentional manipulation: "Target-mask similarity" indicates that the featural and/or meaningfulness relationship between targets and masks was varied; "Attentional set" indicates that observers were either cued to the identity of targets (Eriksen & Collins, 1969), told to attend or not to attend to a mask (Hines & Smith, 1977); "Meaningfulness" indicates that the saliency of target or mask identity was manipulated; "Grouping" indicates that observers were instructed to mentally group a set of stimuli; "---" indicates no manipulation of that type was made in the experiment. The text in the fourth and fifth columns indicates whether a spatial manipulation of attention influenced the level of interruption masking: "Yes" indicates that there was an effect of attention; "No" indicates that there was no effect of attention; "---" indicates that no manipulation of that type was made in the experiment.

Author(s)	Type of Attentional Manipulation		Influence on Masking	
	Target	Mask	Attention to Target	Attention to Mask
Cheal & Gregory (1997)		Visual Search		Yes
Di Lollo, Enns, & Rensink	Visual Search /		No	
(in press)	Location Cueing			
Enns & Di Lollo (1997)	Visual Search		No	
Luck, Hillyard, Mouloua, &		Visual Search		No
Hawkins (1996)				
Shiu & Pashler (1994)		Visual Search		Yes
Tata, Giaschi, & Di Lollo	Visual Search /		No	
(1998)	Location Cueing			
Tata & Giaschi (1999)		Visual Search		Yes

**Table A5.** Influence of spatial manipulations of attention on integration masking. The text in the second and third columns indicate whether attention was varied to the target or to the mask, and the type of attentional manipulation: "Visual search" indicates a target or mask was embedded amongst a variable number of distractors; "Location Cueing" indicates that the location of a target or mask embedded amongst distractors was indicated before onset of a search display; "---" indicates no manipulation of that type was made in the experiment. The text in the fourth and fifth columns indicates whether a spatial manipulation of attention influenced the level of integration masking: "Yes" indicates that there was an effect of attention; "No" indicates that there was no effect of attention; "---" indicates that no manipulation of that type was made in the experiment.

Author(s)	Type of Attention	Influence on Masking		
	Target	Mask	Attention to Target	Attention to Mask
Brehaut, Enns, & Di Lollo (1999)	Attentional blink		No	
Giesbrecht & Di Lollo (1998)	Attentional blink		No	
Visser & Enns (submitted)	Attentional blink		Yes	

**Table A6.** Influence of temporal manipulations of attention on integration masking. The text in the second and third columns indicate whether attention was varied to the target or to the mask, and the type of attentional manipulation: "Attentional blink" indicates the attentional blink paradigm (described in text) was used to vary attention; "unmasking" indicates an analagous paradigm that uses three target stimuli presented in rapid succession (described in text); "---" indicates no manipulation of that type was made in the experiment. The text in the fourth and fifth columns indicates whether a temporal manipulation of attention influenced the level of integration masking: "Yes" indicates that there was an effect of attention; "No" indicates that there was no effect of attention; "---" indicates that no manipulation of that type was made in the experiment.

Author(s)	or(s) Type of Attentional Manipulation		Influence on Masking	
	Target	Mask	Attention to Target	Attention to Mask
Bachmann & Allik (1976)	Target-mask similarity		Yes	
Jacobson (1971)	Target-mask similarity		Yes	
Schiller (1966)	Target-mast similarity		Yes	
Shelley-Tremblay & Mack (1999)	Meaningfulness	Meaningfulness	Yes	Yes

**Table A7.** Influence of non-spatiotemporal manipulations of attention on integration masking. The text in the second and third columns indicate whether attention was varied to the target or to the mask, and the type of attentional manipulation: "Target-mask similarity" indicates that the featural and/or meaningfulness relationship between targets and masks was varied; "Meaningfulness" indicates that the saliency of target or mask identity was manipulated; "---" indicates no manipulation of that type was made in the experiment. The text in the fourth and fifth columns indicates whether a spatial manipulation of attention influenced the level of integration masking: "Yes" indicates that there was an effect of attention; "No" indicates that there was no effect of attention; "---" indicates that no manipulation of that type was made in the experiment.

.

Author(s)	Type of Attentional Manipulation		Influence on Masking	
	Target	Mask	Attention to Target	Attention to Mask
Nazir (1992)	Location Cueing		No	
Reinitz (1990)	Location Cueing		Yes	

**Table A8.** Influence of spatial manipulations of attention on lateral masking. The text in the second and third columns indicate whether attention was varied to the target or to the mask, and the type of attentional manipulation: "Visual search" indicates a target or mask was embedded amongst a variable number of distractors; "Location Cueing" indicates that the location of a target or mask embedded amongst distractors was indicated before onset of a search display; "---" indicates no manipulation of that type was made in the experiment. The text in the fourth and fifth columns indicates whether a spatial manipulation of attention influenced the level of lateral masking: "Yes" indicates that there was an effect of attention; "No" indicates that there was no effect of attention; "---" indicates that no manipulation of that type was made in the experiment.

.

Author(s)	Type of Attentional Manipulation		Influence on Masking	
	Target	Mask	Attention to Target	Attention to Mask
Banks, Larson, & Prinzmetal (1979)	Grouping		Yes	
Banks & White (1984)	Grouping		Yes	
Harcum & Shaw (1974)	Grouping		Yes	
Huckauf, Heller, & Nazir (1999)	Target-mask similarity		No	
Nazir (1992)	Target-mask similarity		Yes	
Wolford & Chambers (1983)	Grouping		Yes	

**Table A9.** Influence of non-spatiotemporal manipulations of attention on lateral masking. The text in the second and third columns indicate whether attention was varied to the target or to the mask, and the type of attentional manipulation: "Target-mask similarity" indicates that the featural and/or meaningfulness relationship between targets and masks was varied; "Grouping" indicates that observers were instructed to mentally group a set of stimuli; "---" indicates no manipulation of that type was made in the experiment. The text in the fourth and fifth columns indicates whether a spatial manipulation of attention influenced the level of lateral masking: "Yes" indicates that there was an effect of attention; "No" indicates that there was no effect of attention; "---" indicates that no manipulation of that type was made in the experiment.