Available online at www.sciencedirect.com

SCIENCE DIRECT®

Cognitive Psychology 48 (2004) 95–126

Cognitive Psychology

www.elsevier.com/locate/cogpsych

Evidence against a central bottleneck during the attentional blink: Multiple channels for configural and featural processing☆

Edward Awh,[a,*] John Serences,[b] Paul Laurey,[a] Harpreet Dhaliwal,[a] Thomas van der Jagt,[a] and Paul Dassonville[a]

[a] *Department of Psychology, 1227 University of Oregon, Eugene, OR 97403-1227, USA*
[b] *Johns Hopkins University, Baltimore, MD, USA*

Accepted 27 May 2003

Abstract

When a visual target is identified, there is a period of several hundred milliseconds when the processing of subsequent targets is impaired, a phenomenon labeled the attentional blink (AB). The emerging consensus is that the identification of a visual target temporarily occupies a limited attentional resource that is essential for all visual perception. The present results challenge this view. With the same digit discrimination task that impaired subsequent letter discrimination for several hundred milliseconds, we found no disruption of subsequent face discrimination. These results suggest that all stimuli do not compete for access to a single resource for visual perception. We propose a multi-channel account of interference in the AB paradigm.

1. Introduction

The attentional blink (AB) refers to a robust limitation in our ability to process sequentially presented target stimuli. When observers attempt to identify two targets

☆ This research was supported by National Institute of Mental Health Grant 1 R01 MH64119 to E.A. We are grateful to Marvin Chun, Vincent Di Lollo, Howard Egeth, and Edward Vogel for helpful discussion and comments.
* Corresponding author.
E-mail address: awh@darkwing.uoregon.edu (E. Awh).

0010-0285/$ - see front matter
doi:10.1016/S0010-0285(03)00116-6

in rapid succession, there is a period of several hundred millisecond after the presentation of the first target (T1) when the accurate identification (or even detection) of the second target (T2) is impaired. The AB has been demonstrated with a wide range of stimuli, including letters, numbers, words, geometric shapes, and colors (Broadbent & Broadbent, 1987; Chun & Potter, 1995; Luck, Vogel, & Shapiro, 1996; Raymond, Shapiro, & Arnell, 1992; Ross & Jolicoeur, 1999; Shapiro, Arnell, & Drake, 1991). In fact, even so-called "preattentive" features (i.e., targets defined by a unique orientation) are misreported when they are presented during the AB period (Joseph, Chun, & Nakayama, 1996). The consistency of this effect suggests interference with a process that is at the core of conscious visual perception. Thus, the models that have been proposed to explain AB interference have consistently invoked capacity-limited processes that are required for the conscious discrimination of all visual stimuli. For example, one class of model suggests that there may be severe capacity limits in the formation of a durable trace in working memory (e.g., Chun & Potter, 1995; Jolicoeur, 1999). By this logic, the processing of T1 occupies a limited-capacity process for consolidation in working memory, preventing the successful consolidation of subsequently presented stimuli for several hundred milliseconds. When T2 is presented during this period, the resulting perceptual trace cannot be consolidated and is therefore vulnerable to overwriting by subsequent stimuli (Giesbrecht & Di Lollo, 1998). A different account offered by Duncan, Ward, and Shapiro (1994) suggests that there is a limited resource for maintaining object representations in a state that is capable of guiding overt behavior. By this view, AB interference results because of competition between different objects for this attentional resource. Thus, T2 is missed because T1 has already occupied the resources that would be necessary for the overt report of the target information. Finally, another account of AB interference (Shapiro, Caldwell, & Sorensen, 1997) suggests that the products of early perceptual analysis are entered into visual short-term memory (VSTM), where each item is assigned a weight that determines its probability of being reported. Because these weights are drawn from a limited resource pool, items coming after the first target may fail to be reported because they are not assigned adequate weights for successful retrieval. While there are important differences between these models, a common feature of these accounts is that they each propose a single resource that is necessary for the perception of all visual stimuli (i.e., for the consolidation, maintenance or retrieval of the relevant target information).

This single-resource view is the key issue at stake in the present research. Our interpretation of this hypothesis motivates the following empirical prediction: AB interference should be observed anytime that the following three conditions are met: (1) The T1 task must demand sufficient resources to deny subsequent stimuli access to stage two processing. (2) The T2 task must require sufficient processing resources to show the effects of the T1 processing load. (3) The first and second targets must be adequately masked (Giesbrecht & Di Lollo, 1998; Grandison, Ghirardelli, & Egeth, 1997; Sieffert & Di Lollo, 1997). Using a procedure that fulfills each of these requirements, we observed a complete absence of AB interference for faces that were presented during the AB period. This result was obtained with the same T1 task that generated long-lasting AB interference for a T2 letter discrimination task (requirement 1). The faces were more difficult to discriminate than the letters that did suffer from AB interference (requirement 2). And the effectiveness of the masks was verified in multiple experiments (requirement 3). These results therefore cast doubt on the idea that a single obligatory process is required for the discrimination of all classes of visual stimuli.

Most previous demonstrations of the AB effect have employed a rapid serial visual presentation (RSVP) technique. In this procedure, observers are asked to discriminate two targets that are embedded within a sequential stream of distractor stimuli. The targets are separated by varying stimulus onset asynchronies (SOAs), and the typical result is that the discrimination or detection of the second target is impaired for several hundred milliseconds after observers have identified the first target. Multiple sources of interference may be present in the RSVP procedure. In addition to the resources that are necessary for discriminating the targets, observers must be able to select only the relevant targets from amongst many irrelevant distractor stimuli, with temporal uncertainty for the onset of the first target. However, some models suggest that AB interference is caused by capacity limitations inherent in the simple act of perceiving and reporting a visual stimulus (e.g., the consolidation of the relevant information in working memory). If this is the case, then it should not be necessary to present irrelevant objects (beyond those needed to mask the targets) or introduce temporal uncertainty in the onset of T1. Duncan et al. (1994) demonstrated this point. They used a simple two-target procedure, in which T1 and T2 were presented at SOAs varying between 0 and 900 ms. When observers were required to discriminate both visual targets, the discrimination of the second target was impaired for several hundred milliseconds after the presentation of the first target. They suggested that the target objects were competing for a limited attentional resource that is necessary for keeping the relevant object information available to influence behavior. We used a procedure similar to that of Duncan et al. (1994) to test the nature of this limitation in sequential target processing.

2. Experiment 1

The purpose of Experiment 1 was to establish a procedure that produces long-lasting AB interference. Current models would suggest that in such a procedure, T2 processing is impaired because T1 processing occupies a resource that is necessary for the perception of any class of visual stimuli. By this account, the same T1 processing load should generate AB interference for any T2 task of equal or greater difficulty.

2.1. Method

2.1.1. Observers

Twelve students from the University of Oregon, between the ages of 18 and 30 years, with normal or corrected-to-normal vision participated in this study for course credit. Six observers participated in the experimental condition and six observers participated in the control condition.

2.1.2. Procedure

Observers were seated 50 cm from the display. They were instructed to maintain fixation throughout each trial. The sequence of events in a single trial of the experimental condition (depicted in Fig. 1) were as follows: (1) A central fixation point and two additional dot markers at the potential target locations (located 1 degree above and below fixation) appeared for 1529 ms. (2) The first target appeared superimposed on the top or bottom marker (location was randomly selected for each trial) for 71 ms. The first target was one of the digits "1," "2," or "3." It was 1° in height and .5° in width. (3) A pattern mask (randomly selected from one of three possible masks) occluded the location of the first target for 71 ms. (4) There were 10 possible stimulus onset asynchronies (SOAs) separating the first and second targets: 0, 59, 118, 176, 235, 294, 353, 412, 529, and 706 ms. SOA was randomly intermixed between trials. (5) At the selected SOA, the second target was presented either to the right or the left of the fixation point for 71 ms. All of the second targets fit exactly inside an imaginary rectangle with a height of 5.5° and a width of 4°. The distance from the center of this rectangle to fixation was 4.5°. The second target was the upper case version of one of the letters "J," "K," or "L." (6) The second target was completely occluded by a mask for 59 ms (see Fig. 1 for examples of masking stimuli). There were three versions of each mask type, one of which was randomly selected during each trial. (7) After both targets had been presented, observers reported the identity of both targets with unspeeded key presses. The sequence of events in the control condition was identical to that of the experimental condition, except that observed were required to report only the identity of the second target.

2.1.3. Design

Observers participated in 15 blocks of 30 trials in the experimental and control conditions, yielding 45 observations at each SOA in each condition. One block of

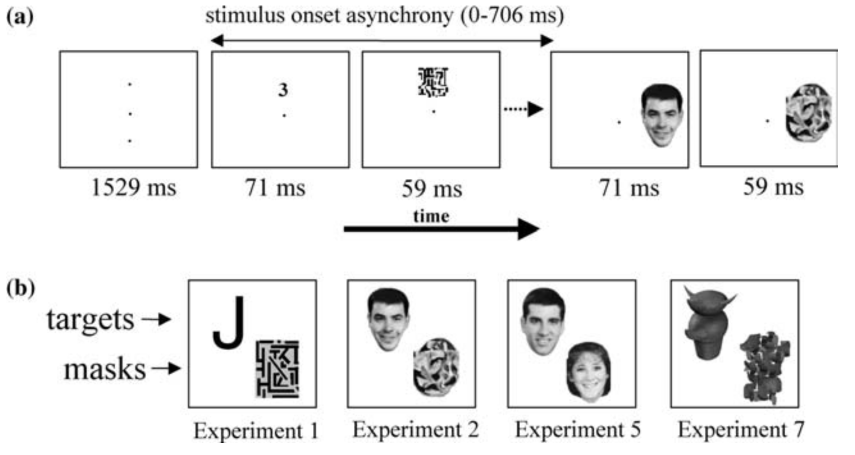

Fig. 1. Trial schematic: (a) The sequence of events in a single trial of Experiment 2. (b) Examples of the second target stimuli and masks used in Experiments 1, 2, 5, and 7.

practice was administered before each condition. In this experiment and all those that follow, the data from trials with inaccurate responses to the first target were excluded from analysis, but the same pattern of results was obtained when these trials were included.

2.2. Results and discussion

The mean accuracy for the first target was 84%. The primary goal of the first experiment was to establish that the procedure we used would replicate previous observations of the attentional blink. We measured target two accuracy as a function of SOA in an *experimental* condition that required the report of both the first and second target, and in another *control* condition in which the first target was ignored. The control condition enabled us to account for the effects of simple perceptual interference caused by the stimulus display. This experiment confirmed previous observations that this simple two-target procedure can produce robust AB interference (Duncan et al., 1994; Ward, Duncan, & Shapiro, 1997). Fig. 2 illustrates how the identification of the second target was impaired in the experimental relative to the control condition ($F[1, 10] = 12.4$, $p < .01$). Replicating previous studies, we observed an impairment that lasted for several hundred milliseconds, with accuracy in the experimental and control conditions gradually converging as SOA increased, ($F[9, 90] = 2.4$, $p < .02$).

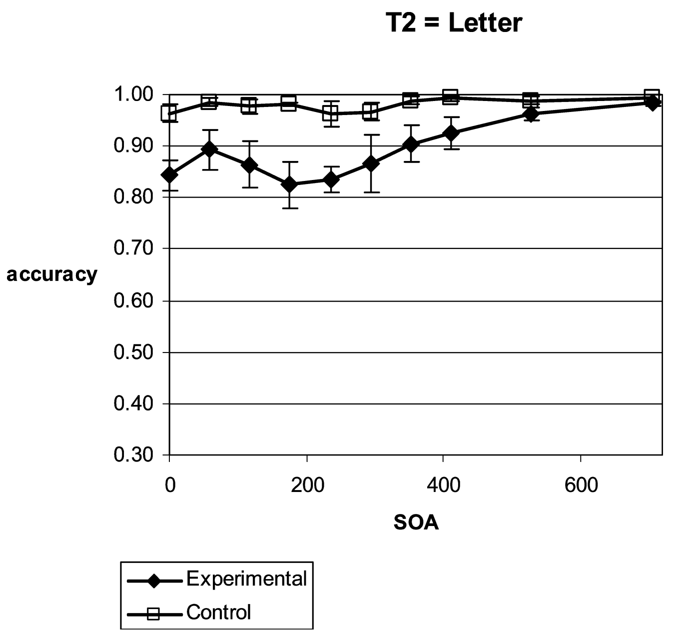

Fig. 2. Mean accuracy of second target (T2) report, given correct report of the first target (T1) as a function of stimulus onset asynchrony (SOA) and experimental condition in Experiment 1.

3. Experiment 2

Experiment 1 established that this version of the two-target procedure produces AB interference that lasts over 500 ms. Experiment 2 provides a direct test of the hypothesis that this interference reflects the disruption of a process that is central to the discrimination of any visual stimulus. The T1 processing load was identical in Experiment 2, but the T2 stimuli were faces instead of letters.

3.1. Methods

3.1.1. Observers

Eight students from the University of Oregon with normal or corrected-to-normal vision received course credit for their participation.

3.1.2. Procedure

The sequence of events in each trials is depicted in Fig. 1. All methodological details were identical to those of Experiment 1, except that the second target was a photograph of a face instead of a letter. The potential T2 stimuli and masks are illustrated in Fig. 3. Observers discriminated between three possible face targets (all unknown to the observers). One of three possible masks (created by scrambling parts of the face stimuli) was randomly selected for each trial. Observers reported the identity of the face that was presented in each trial by pressing either the "A," "S," or "D" key. Each of these keys was associated with a single face by means of a diagram that showed each face with an arrow pointing to the proper response key. This diagram was available to observers throughout the experiment. Each observer participated in both the experimental and the control conditions. The order of these conditions was counterbalanced across observers.

3.2. Results and discussion

The mean accuracy for the first target was 94%. When faces were presented in the second target position, there was no trace of an attentional blink. As Fig. 4 illustrates, accuracy was identical in the experimental and control conditions ($F[1,7] = .05,\ p = .84$). This match between experimental and control accuracy was maintained at all SOAs tested; there was no interaction of condition and SOA, ($F[9,63] = .57, p = .82$). Thus, although the task of identifying the first target was identical in Experiments 1 and 2, there was no evidence of an attentional blink when the second targets were faces. We did, however, observe a significant effect of SOA ($F[9,63] = 4.5, p < .01$), reflecting a gradual rise in accuracy from 87 to 91% as SOA increased from 0 to 706 ms.[1]

[1] This effect of SOA is consistent with the possibility that AB interference was occurring in both the experimental and control conditions of this experiment—if observers had inadvertently processed the T1 stimulus in the control condition. However, Experiment 3 replicates the results of the first two experiments, while ensuring equivalent T1 processing in the face and letter conditions.

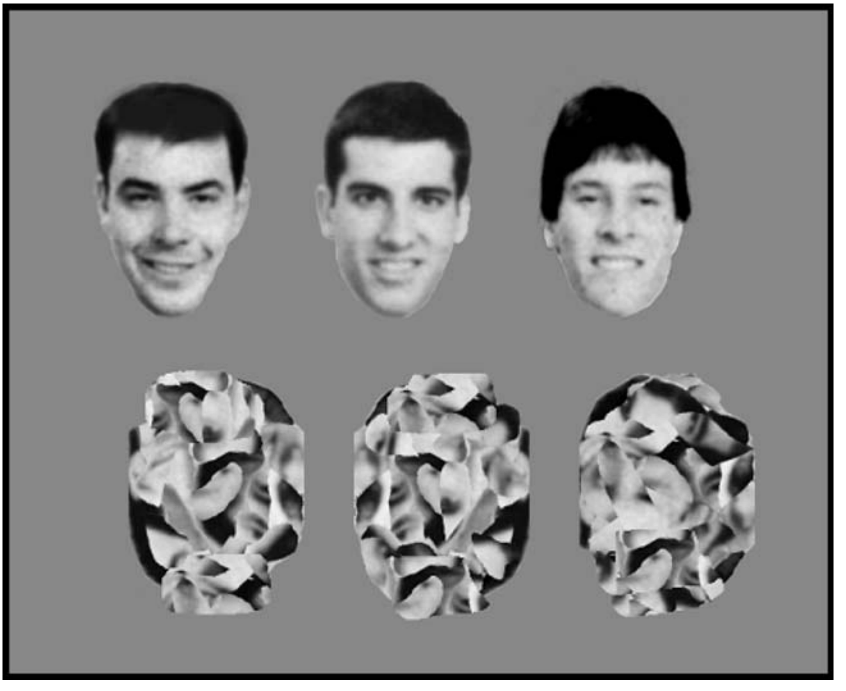

Fig. 3. The T2 faces and masks employed in Experiment 2. Individual targets and masks were randomly paired during each trial.

A previous study has shown that one's own name can be detected from a stream of distractor words during the attentional blink period (Shapiro et al., 1997). This experiment made the important point that the words with high semantic salience—in this case due to very high familiarity—are more easily detected during the AB period than other words. This result provides strong support for a late selection account of this phenomenon, because the probability of successful discrimination was determined by a semantic feature of the stimuli. However, the results with personal names do not rule out the view that conscious perception of all classes of visual stimuli is impaired during the attentional blink. Instead, they show that specific members of the word class are less likely to show the effects of AB interference, perhaps because the threshold of activation needed to detect these stimuli is significantly lower than for other words (Shapiro et al., 1997). Our results make a qualitatively different point. The present experiment required observers to discriminate between three completely novel face stimuli. These targets could not have enjoyed the same benefits that Shapiro et al. documented for personal names (i.e., lower thresholds for activation induced by massive amounts of experience). Nevertheless, Experiment 3 provides a direct test of the idea that this face discrimination task was simply easier than the letter task used in Experiment 1.

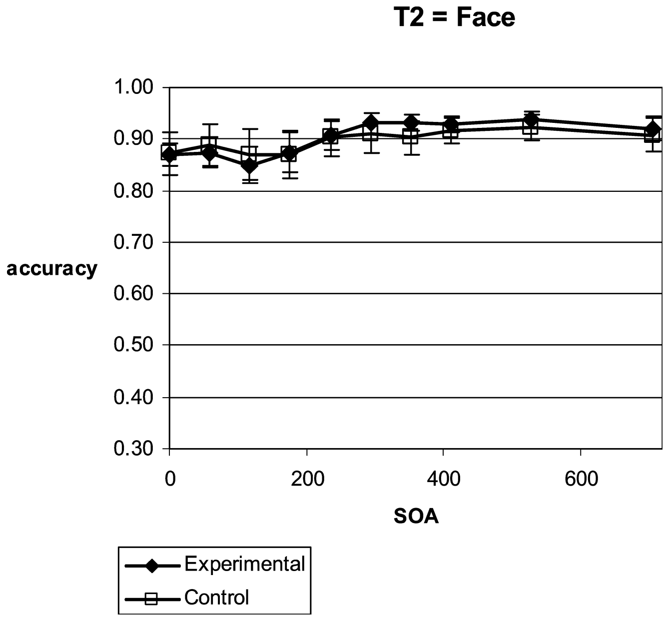

Fig. 4. Mean accuracy of second target (T2) report, given correct report of the first target (T1) as a function of stimulus onset asynchrony (SOA) and experimental condition in Experiment 2.

4. Experiment 3

Experiments 1 and 2 suggest a clear difference in the processing of faces and letters during the AB period. While letter discrimination showed a long-lasting deficit after the attentive processing of T1, the faces were equally well discriminated in the experimental and control conditions. However, the conclusion that faces are resistant to AB interference depends upon the assumption that T1 processing was identical in Experiments 1 and 2. An alternative possibility is that observers devoted more attention to T1 processing in the letter condition than in the face condition. If so, then T1 processing in Experiment 1 may have absorbed more of the limited capacity process that has been hypothesized to generate AB interference. One problem for this account is the fact that T1 accuracy was actually higher in the context of the face discrimination (94%) than in the context of the letter discrimination (84%) ($t(9) = 3.8$, $p < .01$). This contradicts the idea that observers devoted more attentional resources to T1 processing in Experiment 1. We also considered the related hypothesis that T1 processing was different in the control conditions of Experiments 1 and 2. That is, observers in Experiment 2 may have inadvertently processed the T1 stimulus during the control condition, despite our instructions to ignore it. In this case, there may have been AB interference in both the experimental and control conditions of Experiment 2, preventing us from observing the typical contrast between the conditions. This possibility is consistent with our finding of a significant effect of SOA in both the experimental and control conditions of Experiment 2. Of course, AB interference during the control condition of Experiment 2 would compromise our use of this condition as a baseline for assessing performance in the experimental condition.

Another alternative explanation is that face discrimination did suffer from AB interference, but the task was so easy that we were unable to detect it. Many studies have documented the high level of expertise that characterizes face processing (Gauthier, William, Tarr, & Tanaka, 1998; Tanaka & Gauthier, 1997). Because of the tremendous amount of experience that people have with this important perceptual task, these stimuli can be discriminated more efficiently than other equally homogenous categories of stimuli. Thus, it would be prudent to measure the difficulty of any perceptual task that apparently avoids AB interference.

Experiment 3 addressed all of the above considerations. The procedure was very similar to that of Experiments 1 and 2, except that the T2 stimulus was randomly varied between the same face and letter stimuli that were used in the previous experiments. T2 was a face on half of the trials, and a letter on the other half of the trials. We reasoned that T1 processing would be identical in the face and letter conditions because in the vast majority of trials T1 processing would be well under way by the time the second target appeared. In addition, the overall difficulty of the face and letter tasks was equated by setting the exposure duration of each target category on a within-subject basis. Finally, using this within-subject measure of exposure duration as a dependent measure allowed a direct comparison of the difficulty of the face and letter discrimination tasks.

4.1. Methods

4.1.1. Observers

Sixteen students from the University of Oregon, between the ages of 18 and 30 years, with normal or corrected-to-normal vision participated in this study for course credit.

4.1.2. Procedure

Each observer participated in 14 blocks of 24 trials in the experimental condition and 14 blocks of 24 trials in the control condition. The order of these conditions was counterbalanced across observers. One block of practice was administered before each condition.

4.1.3. Stimulus presentation

Within each trial of the experimental and the control conditions, the T2 object had a 50% chance of being either a face or a letter. Four SOAs were tested: 0, 118, 235, and 706 ms. For each observer, the exposure duration of each stimulus category was adjusted by means of a 120 trial timing procedure. The stimulus was initially presented at an exposure duration of 72 ms (six monitor refresh cycles at 85 Hz) in the zero SOA condition of the control task. If the stimulus was accurately identified, the exposure duration was reduced by one monitor refresh cycle. If the stimulus was inaccurately identified, the exposure duration was increased by two monitor refresh cycles. After 120 trials, the average exposure duration over the last 30 trials was calculated and used for the remainder of the experimental trials. This procedure was successful at producing about 70% accuracy for each stimulus category in the control condition (Fig. 6).

4.2. Results and discussion

Fig. 5 illustrates T1 accuracy as a function of the T2 stimulus and SOA. At all SOAs tested, T1 accuracy was identical in the face and letter conditions. Thus, a two-way ANOVA of T1 accuracy with T2 stimulus (face or letter) and SOA: (0, 118, 235, or 706 ms) revealed no main effect of task ($F[1, 15] = .03$, $p = .87$) and no interaction of task and SOA ($F[3, 45] = 2.1$, $p = .112$). We were confident, therefore, that the T1 processing load was equated in the face and letter conditions. This was expected, because T1 processing was usually well under way before the T2 stimulus had even appeared. For the same reason, this design also ruled out the possibility that T1 was inadvertently processed in the control condition of only the face condition.

Fig. 6 shows T2 accuracy as a function of task (face or letter), condition (experimental or control) and SOA. As the graph illustrates, the discrimination of T2 letters was significantly impaired in the experimental condition relative to the control condition. But no significant difference between experimental and control conditions was observed with T2 faces. A three-way ANOVA of T2 accuracy with T2 stimulus-type, condition, and SOA as factors confirmed these observations. There was a main effect of condition ($F[1, 15] = 4.9$, $p < .05$), with worse performance in the experimental than in the control condition. A main effect of SOA ($F[3, 45] = 12.9$, $p < .01$) reflected a slight increase in accuracy as SOA increased. A significant interaction of T2 stimulus-type and SOA ($F[3, 45] = 4.8$, $p < .01$) reflected a larger

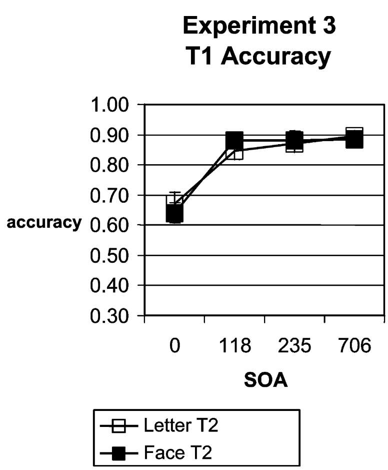

Fig. 5. Mean accuracy of first target (T1) report as a function of stimulus onset asynchrony (SOA) and the category of the second target (T2) in Experiment 3.

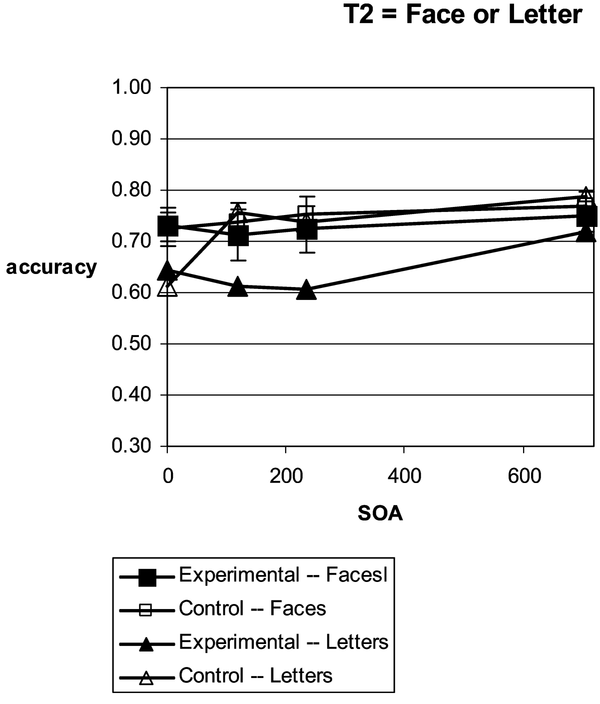

Fig. 6. Mean accuracy of second target (T2), given correct report of the first target (T1) as a function of stimulus onset asynchrony (SOA), the category of the second target stimulus and experimental condition in Experiment 3.

change in accuracy as SOA increased in the letter task than in the face task. An interaction between condition and SOA ($F[3,45] = 8.5$, $p < .01$) showed that the difference between the experimental and control conditions was restricted to the second and third SOAs. Finally, the most telling result was a three-way interaction of T2 stimulus-type, condition, and SOA ($F[3,45] = 2.8$, $p < .05$), showing that the difference between the experimental and control conditions (apparent in the second and third SOAs tested) was restricted to the letter condition. These data confirmed the resistance of faces to AB interference while ensuring that T1 processing was equivalent in the face and letter conditions.

Experiments 1 and 2 left open the possibility that we were unable to detect AB interference with the faces, because of unanticipated interference in the control condition—perhaps due to unnecessary processing of the T1 stimulus in the face control condition. In addition, we considered the possibility that fewer resources were dedicated to the T1 task in the experimental condition of Experiment 2 (with faces). However, the results of Experiment 3 show the same interaction between T1 processing load and stimulus class, while ensuring the equivalence of T1 processing load for the faces and letters. Thus, the absence of AB interference for faces does not appear to derive from differential processing of the first target.

Recall that the exposure duration for the face and letter stimuli was set on a within-subject basis in Experiment 3. The procedure was successful in achieving approximately equivalent accuracy for the face and letter discriminations in the control condition. Moreover, the resulting exposure times suggested that the faces were more difficult to perceive than the letters. The mean exposure duration for the faces (63 ms) was significantly higher than the mean exposure duration for the letters (44 ms) ($t(13) = 3.1$, $p < .01$). These data suggest that the lack of AB interference for the faces was not a simple result of an easy discrimination task. There are many potential sources of difficulty, however. One possibility is that the exposure times for the face stimuli were influenced by perceptual factors that are unrelated to the root source of AB interference. Multiple studies have suggested that AB interference affects a postperceptual stage of processing (e.g., Jolicoeur & Dell' Acqua, 2000; Shapiro et al., 1997; Vogel, Luck, & Shapiro, 1998) such as the consolidation of information in working memory. This leaves open the possibility that the faces, while taking longer to perceive overall, place little strain on the specific postperceptual process that leads to AB interference. Experiment 4 tested this hypothesis by reversing the order of the target events in Experiment 2. When observers in Experiment 2 discriminated the digit before the face, no AB interference was observed. If this null effect came about because face discrimination places little demand on AB-related processes, then the same result should be obtained when the order of the targets is reversed. Specifically, if the face task does not demand the resources that are critical to inducing AB interference, then there would be no reason to predict that these resources should be unavailable for processing the subsequent digit target.

5. Experiment 4

5.1. Methods

5.1.1. Observers

Ten students from the University of Oregon between the ages of 18 and 30 years, with normal or corrected-to-normal vision, participated in this study for course credit.

All methodological details were identical to those of Experiment 2, except that the order of the stimuli was reversed. The faces (positioned just as they were in Experiment 2) appeared first and were followed by the digit stimuli using the same range of SOAs tested in Experiments 1 and 2.

6. Results and discussion

The mean accuracy for the first target was 85%. Fig. 7 illustrates T2 accuracy as a function of condition (experimental or control) and SOA. It is clear from this graph that there was a significant disruption of performance in the experimental condition

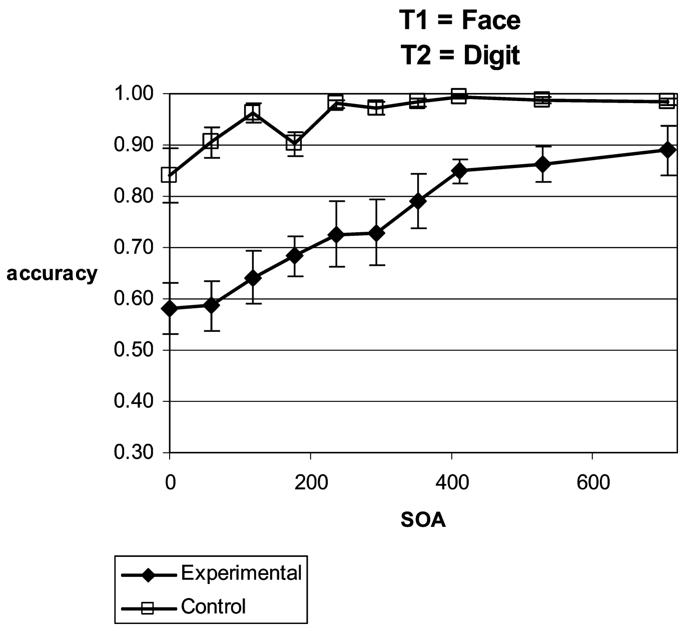

Fig. 7. Mean accuracy of second target (T2) report, given correct report of the first target (T1) as a function of stimulus onset asynchrony (SOA) and experimental condition in Experiment 4.

relative to the control condition. An ANOVA with subjects, condition (experimental vs. control), and SOA as factors showed a main effect of condition ($F[1,9] = 37.7$, $p < .01$), reflecting much higher accuracy overall in the control condition (94.6%) than in the experimental condition (70.4%). There was also a significant interaction of condition and SOA ($F[9,81] = 4.8$, $p < .01$) that resulted from the gradual recovery in the experimental condition as SOA increased. Finally, there was a main effect of SOA ($F[9.81] = 20.4$, $p < .01$) that resulted from higher accuracy as SOA increased.

Reversing the order of the stimuli had a dramatic impact on performance in the experimental condition. While T2 processing was identical in the experimental and control conditions when the digits were processed before the faces, there was a severe disruption in T2 processing when the faces were processed before the digits. This result is inconsistent with the hypothesis that face discrimination places relatively low demands on the processes that cause AB interference. Instead, we found that even though face discrimination recruits a process that generates strong AB interference for digits, digit discrimination does not occupy a critical resource for the discrimination of faces. This asymmetry highlights the possibility that AB interference does not result from the disruption of a single resource for visual perception. If the disruption of T2 accuracy in Experiment 4 resulted from the heavy recruitment of a single obligatory resource, then the resource-demanding face task should have been disrupted in Experiment 2 when it was in competition with the digit discrimination task. By contrast, the asymmetry might be explained if observers had access to an additional processing channel for the discrimination of the face stimuli. This hypothesis is further developed in Section 13.

The combined results of Experiments 2 and 4 make an important point about top-down control in this two-target paradigm. The visual displays in Experiments 2 and 4 were identical when the SOA between T1 and T2 was 0 ms (i.e., T1 and T2 were presented simultaneously). Nevertheless, accuracy in these two conditions was remarkably different. Consider observers' performance in the experimental condition when the targets were presented simultaneously. In Experiment 2, the digits and faces were identified with a mean accuracy of 83 and 87%, respectively. However, with an identical display in Experiment 4 (at the 0 SOA) the digits and faces were discriminated with a mean accuracy of 58 and 81%, respectively. Unpaired *t* tests show that both face discrimination ($t(16) = 20.0$, $p < .01$) and digit discrimination ($t(16) = 24.8$, $p < .01$) were significantly impaired in Experiment 4 relative to Experiment 2. The target displays in these trials were identical. This suggests that accuracy was affected by differences in the observers' top-down attentional settings. Of course, the simple decision to direct attention to the T1 object is a necessary precondition for AB interference. But these data suggest that the right top-down settings can prevent interference even when both T1 and T2 are fully processed. One possibility is that the observers were influenced by the expected order of the target stimuli. In Experiment 2, the digits appeared first in 90% of the trials. In Experiment 4, the faces appeared first 90% of the time. Di Lollo, Kawahara, Zuvic, and Visser (2001) have suggested that the visual system may be reconfigured on a moment-to-moment basis in order to prepare for the demands of a specific perceptual task. It may be that observers configured their visual systems to process digits in Experiment 2 and faces in Experiment 4. The results suggest that there may be asymmetric effects of preparing to process face and digit stimuli. When the visual system is configured to process digits, these settings are still compatible with the accurate discrimination of subsequent face stimuli. However, when the visual system is configured to process faces, there is strong interference between face and digit processing. The implications of this asymmetry are discussed further in Section 13. Although the overt discrimination of the T1 stimulus is necessary for observing AB interference, the T1 processing load does not always determine whether interference is observed. Top-down control settings can prevent or induce interference when the overall processing requirements are held constant.

7. Experiment 5

Previous studies of the attentional blink have shown that if the second target is not masked effectively, the attentional blink is eliminated (Brehaut, Enns, & Di Lollo, 1999; Giesbrecht & Di Lollo, 1998). The interruption masks we used in Experiment 2 were clearly sufficient to eliminate visual persistence, but we have considered a more subtle possibility. It has been suggested that a key component of the attentional blink is the substitution of the second target representation by the interruption mask as the object for eventual conscious registration (Giesbrecht

& Di Lollo, 1998). This raises the possibility that the absence of an attentional blink in Experiment 2 derived from the use of masks that were ineffective substitutes for the faces, due to non-overlapping stages of processing for the two objects. This hypothesis is made plausible by suggestions that face processing may be mediated by a unique neural substrate (Kanwisher, McDermott, & Chun, 1997).

Another potential explanation for the preserved processing of faces in these experiments has to do with the relative salience of face stimuli. Some studies have suggested that faces may be more effective at capturing attentional resources than other classes of stimuli (e.g., Ro, Russell, & Lavie, 2001). Perhaps the resistance of faces to AB interference is a result of a natural advantage in the competition for visual attention. One problem with this account is the fact that T1 digit processing was no worse when faces were presented in the T2 position than when letters were presented in the T2 position. If the faces were capturing attentional resources that were necessary for T1 processing, then presumably there would have been some effect on T1 accuracy. Nevertheless, it would be informative to observe whether faces can avoid AB interference even when they could be overwritten by another equally salient stimulus. Experiment 5 replicated Experiment 2 using another set of face stimuli as masks for the second target. The processing of these face masks should have maximal overlap with that of the face targets, providing optimal conditions for object substitution. In addition, because the masks were drawn from the same stimulus class as the faces, we could assess performance while the faces were in direct competition with equally salient distractor objects.

7.1. Method

Experiment 5 was an exact replication of Experiment 2, except that during each trial the masking stimulus was randomly selected from one of the three female faces (illustrated in Fig. 8). Eight students from the University of Oregon received course credit for their participation.

Fig. 8. The three female faces that served as masks in Experiment 5. Individual targets and masks were randomly paired during each trial.

7.2. Results and discussion

The mean accuracy for the first target was 93%. As Fig. 9 illustrates, the resistance of faces to AB interference was replicated under the new masking conditions; accuracy in the experimental and control conditions was identical, ($F[1,7] = .03$, $p = .87$). The unimpaired performance in the experimental condition was maintained throughout all tested SOAs; there was no interaction of condition and SOA ($F[9,63] = .59, p = .80$). Although performance in the experimental and control conditions remained identical with the face masks, there was a strong main effect of mask type on second target performance. The average accuracy for target two discrimination dropped substantially from 91% in Experiment 2 to 64% in Experiment 5, $t(14) = 5.2$, $p < .01$. This 27% drop in accuracy suggests that the face masks were a powerful source of competition for the face targets. Consistent with this claim, the relative salience of the targets and masks should have been similar in this study, because both sets of faces were novel to the observers. Thus, this experiment bolsters the evidence from Experiment 3 that the overall difficulty of the face discrimination cannot explain the absence of AB interference for these stimuli. More importantly, Experiment 5 suggests that our initial observations of accurate face discrimination during the AB period did not result from the use of masks that were processed within an independent perceptual stream.

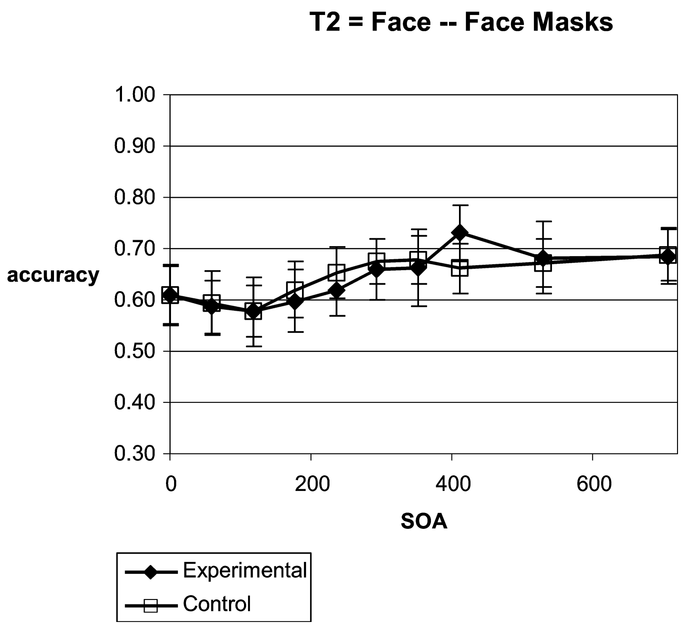

Fig. 9. Mean accuracy of second target (T2) report, given correct report of the first target (T1) as a function of stimulus onset asynchrony (SOA) and experimental condition in Experiment 5.

There was a significant main effect of SOA in Experiment 5 ($F[1,7] = 7.0$, $p < .01$), reflecting lower accuracy (in both the experimental and control conditions) at the earliest SOAs (up to about 176 ms after the onset of T1). Just like the main effect of SOA in Experiment 2, this result raises the possibility that observers inadvertently processed the T1 stimulus in the control condition. If this had caused AB interference in the control condition of this experiment, then our sensitivity to AB interference in the experimental condition would be compromised. The same concern was addressed directly by Experiment 3. Experiment 3 demonstrated AB interference for letters but not faces while ensuring the equivalence of T1 processing in each condition. This confirmed that the manipulation of T1 processing load (across the experimental and control conditions) was effective at producing long-lasting AB interference for letters. Thus, the absence of AB interference for faces (in Experiment 3) could not be ascribed to a failed manipulation of T1 processing load. Based on this finding, we argue that the SOA effect in Experiment 5 (which was equivalent in the experimental and control conditions) is most likely a result of perceptual interference caused by the stimulus display rather than AB interference. Indeed, the primary purpose of the control condition was to rule out stimulus-driven sources of interference. Of course, the stimulus conditions were not identical in Experiments 3 and 5. Would observers have been more likely to process T1 in the control condition of Experiment 5? This hypothesis relies on the assumption that increasing the difficulty of the T2 discrimination (by means of a more salient T2 mask) would increase the tendency for observers to process T1 in the control condition. But it seems more likely that increasing the difficulty of the T2 task would have just the opposite effect. Hence, we conclude that faces avoided AB interference in Experiment 5 despite the use of faces as masking stimuli.[2]

8. Experiment 6

So far, we have explored several explanations for the resistance of faces to AB interference. The face advantage cannot be explained by the differential processing of T1 in the face and letter conditions. The perceptual difficulty of the face discrimination does not provide an explanation either; this face task is more difficult than the

[2] We further explored the possibility of AB interference in the experimental condition of Experiment 5 by conducting paired comparisons of the accuracy at early SOAs compared with accuracy at the 706 ms SOA (see Raymond, Shapiro, & Arnell, 1995 for a similar approach). This analysis (uncorrected for multiple comparisons) revealed a significant reduction in accuracy at only the 156 ms SOA ($t(7) = 1.9$, $p < .05$). Reduced accuracy at only the 156 ms SOA is inconsistent with the typical time course of AB interference. For example, a similar analysis of the experimental condition of Experiment 1 (in which the same T1 task was used) revealed significant differences between accuracy at the final SOA and all SOAs from 0 up to 353 ms. This kind of long-lasting interference is more consistent with previous demonstrations of AB interference. Thus, we conclude that the effect of SOA in Experiment 5 is likely to result from perceptual interference (e.g., forward masking from the T1 presentation) that is qualitatively distinct from the attentional blink.

letter task that does show AB interference. The face task is capable of inducing a sizable attentional blink for subsequently presented letters; this rules out the possibility that face discrimination places a low load on the specific subset of processes that lead to AB interference. Finally, even when the faces were masked by other faces—optimizing the chances for object substitution—no AB interference was observed. To summarize, the difficulty of the face discrimination (manipulated in multiple ways) and the size of the T1 processing load cannot explain the difference between face and letter processing in this paradigm.

Our primary conclusion from these results is that the disruption of T2 processing during the AB period may not be a result of competition for a single obligatory resource for visual perception. As an alternative, we propose a multi-channel account of AB interference. By this view, there exist multiple routes by which stimuli can be explicitly discriminated during the AB period. AB interference is therefore a product of the overlap between the processing channels that guide the discrimination of T1 and T2. If T1 processing leaves open a channel that is sufficient to guide T2 processing, then AB interference can be avoided. In the present experiments, face discrimination may have access to a processing channel that is unavailable for letters and digits.

What is the nature of this extra processing channel? For the present purposes, we rely on a distinction elucidated by Farah, Wilson, Drain, and Tanaka (1998) between a "holistic" or "configural" mode of processing and a "feature-based" mode of processing. We assume that the digit discrimination task is guided by the feature-based system, while the face discrimination task recruits both the feature-based and the configural processing systems. There is precedent for the idea that faces evoke multiple representational codes. Gauthier and Tarr (2002) provide evidence that both feature-based and configural representations are activated even for objects known to evoke configural codes. For example, they found cases in which configural processing aided the identification of only a subset of the parts within an object. Thus, while it is clearly established that face discrimination relies on configural cues, the representation of individual parts within these configural objects should also be acknowledged. With these assumptions, all of our observed data can be explained by the multi-channel hypothesis. When T1 processing occupies the feature-based system, T2 faces could still be discriminated based on information from the configural system—consistent with the results of Experiments 2, 3, and 5. However, when T1 is a face then both feature-based and configural processes would be occupied when T2 was presented. This would predict AB interference for subsequent digits, consistent with the results of Experiment 4.

Experiment 6 tested another clear prediction of the multi-channel hypothesis that was not addressed by Experiments 1–5. AB interference should be observed if both T1 and T2 are faces, because T1 would occupy all the perceptual channels that could guide the discrimination of the second face. In addition, a demonstration of AB interference for faces can rule out the possibility that faces are resistant to AB interference simply because they have a unique set of low-level features. For example, the presence of low spatial frequency information and configural information in these stimuli may have rendered them less susceptible to masking—the results of Experiment 5 notwithstanding. Experiment 6 disconfirmed this hypothesis by demonstrating that faces do not always avoid AB interference.

9. Experiment 6

9.1. Methods

9.1.1. Observers

Eight students from the University of Oregon, between the ages of 18 and 30 years, with normal or corrected-to-normal vision participated in this study for course credit.

9.1.2. Procedure

Each observer participated in six blocks of 30 trials in the experimental condition and six blocks of 30 trials in the control condition. The order of these conditions was counterbalanced across observers. One block of practice was administered before each condition.

9.1.3. Stimulus presentation

All aspects of the task were identical to those of Experiment 2, with some changes in the geometry of the stimulus display and a change in the stimuli used as T1. The changes were as follows: (1) The two dots that marked the potential locations of T1 were located 4° above and below fixation. (2) T1 was one of the three female faces that were the same size as the T2 faces (4° wide by 5.5° tall). (3) The same face stimuli used in Experiment 2 served as the T2 stimuli again. The potential locations were to the left and right of fixation, with the center of the image located 6° from fixation. (4) Four SOAs were tested (0, 235, 706, and 1176 ms).

9.2. Results and discussion

The mean accuracy for the first target was 76%. T2 accuracy is displayed in Fig. 10, as a function of condition, and SOA. There was clear evidence of long-lasting AB interference, with worse performance in the experimental condition than in the control condition, and gradual recovery over the four SOAs tested. A two-way ANOVA with subjects, condition (experimental vs. control) and SOA as factors showed a main effect of condition ($F[1,7] = 50.4$, $p < .01$), and an interaction of condition and SOA ($F[3,21] = 5.6$, $p < .01$). There was also a main effect of SOA ($F[3,21] = 79.0$, $p < .01$). This demonstration of AB interference with faces in the T1 and T2 positions is consistent with a multi-channel view, because the processing of the first target should have occupied all of the processing channels that could guide the discrimination of the second target. In addition, Experiment 6 shows that the masks that were used for the face stimuli (in the present experiment and in Experiments 2, 3, and 4) will produce AB interference when there is sufficient overlap in the processing requirements for T1 and T2.

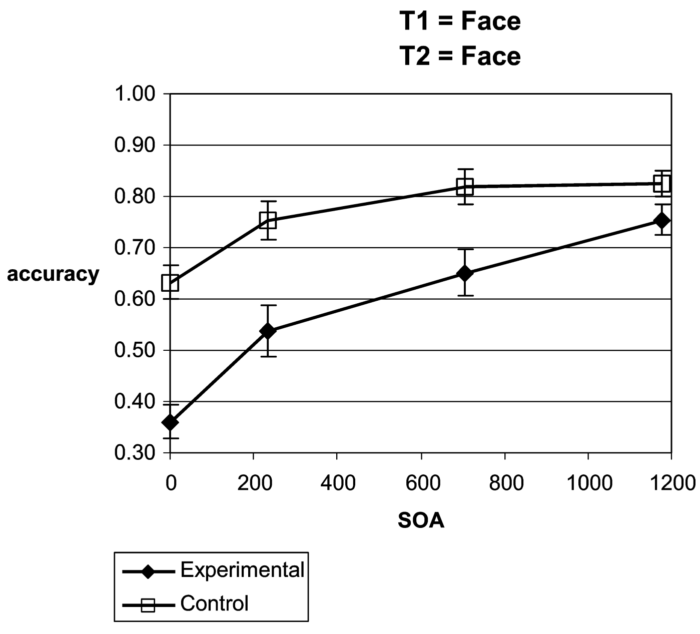

Fig. 10. Mean accuracy of second target (T2) report, given correct report of the first target (T1) as a function of stimulus onset asynchrony (SOA) and experimental condition in Experiment 6.

10. Experiment 7

Much research has focused on the high level of expertise that characterizes face processing (Gauthier et al., 1998; Tanaka & Gauthier, 1997). Thus, the resistance of faces to AB interference may be partially dependent on observers' perceptual expertise with these stimuli. Would a novel class of stimuli that can be discriminated based on configural information prove to be resistant to AB interference? We explored this question in Experiment 7 by employing Greebles (Tanaka & Gauthier, 1997). These stimuli have been used in previous research that explores the development of perceptual expertise for novel stimulus categories. Given adequate experience identifying Greebles, behavioral performance with these stimuli can be shown to mirror specific characteristics of face perception. For example, trained Greeble observers recognize individual Greeble parts more quickly within the context of the original object in which the parts were learned compared to within a new configuration of parts. Greeble experts also show activation of the "fusiform face area" when they view these stimuli; this activity in the inferotemporal cortex has been used as a marker of perceptual expertise in various studies (Tarr & Gauthier, 2000). These results suggest that Greebles contain the necessary configural information for the development of perceptual expertise. Thus, testing untrained observers with these stimuli provided an opportunity to observe whether the resistance of faces to AB interference is contingent upon the use of stimuli that are discriminated through expert processing.

10.1. Methods

10.1.1. Observers

Eight volunteers from the local community were paid for their participation in a 1-h session. The observers were between the ages of 18 and 30 years, with normal or corrected-to-normal vision.

10.1.2. Procedure

Each observer participated in six blocks of 30 trials in the experimental condition and six blocks of 30 trials in the control condition. The order of these conditions was counterbalanced across observers. One block of practice was administered before each condition.

10.1.3. Stimulus presentation

All aspects of the procedure were identical to those of Experiment 2 with the following exceptions: Three Greeble stimuli were used in place of the face stimuli. The three possible T2 stimuli and masks are illustrated in Fig. 11. The Greebles were masked with one of three masks. Observations were collected using six dif-

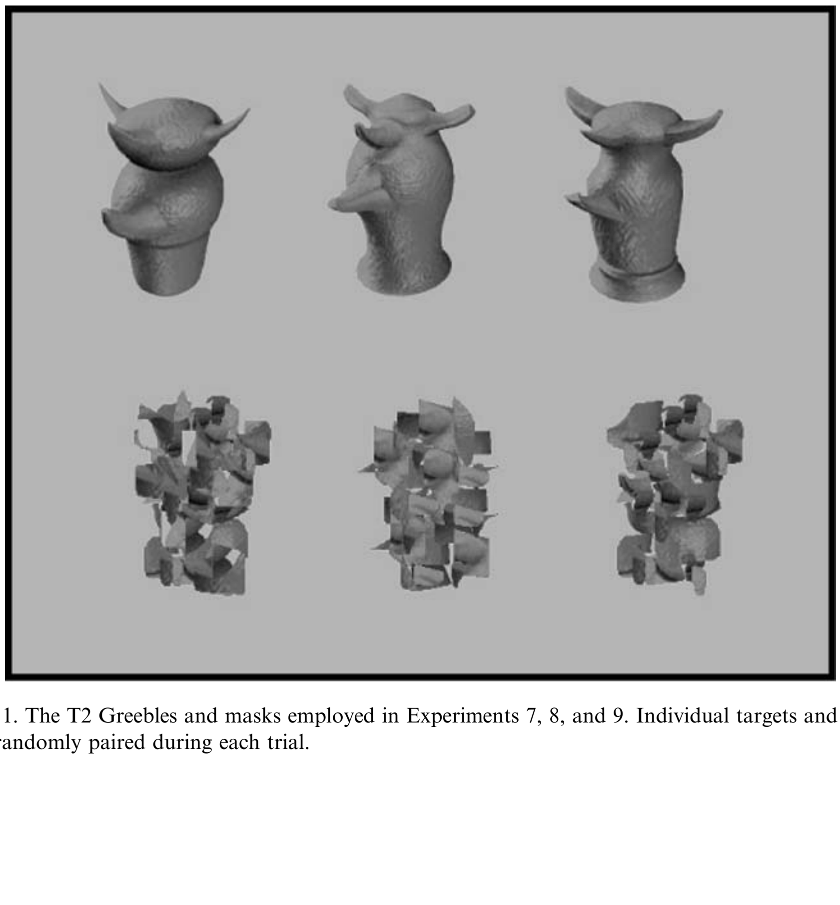

Fig. 11. The T2 Greebles and masks employed in Experiments 7, 8, and 9. Individual targets and masks were randomly paired during each trial.

ferent SOAs: (118, 294, 400, 506, 600, and 706 ms). The exposure duration of each target was set at 94 ms in order to maintain accuracy at about 70% in the control condition. In order to facilitate the observers' learning of the Greeble identities, the target exposure durations were raised to 141 ms in the practice block.

10.2. Results and discussion

The mean accuracy for the first target was 97%. T2 accuracy is displayed in Fig. 12 as a function of condition and SOA. As Fig. 12 illustrates, there was little evidence of impaired performance in experimental condition. A two-way ANOVA with subjects, condition (experimental vs. control), and SOA as factors did not show a main effect of condition ($F[1, 7] = .03$, $p = .87$). There was, however, a significant interaction of SOA and condition ($F[5, 35] = 3.4$, $p < .015$). Paired t tests confirmed that there was a marginally significant impairment in the experimental condition at the 118 ms SOA ($t(7) = 2.3$, $p = .06$). But no other SOA showed any trace of impaired performance in the experimental condition. Finally, there was a significant main effect of SOA, with accuracy gradually improving as SOA increased ($F[5, 35] = 14.2$, $p < .01$). The early recovery from interference in Exper-

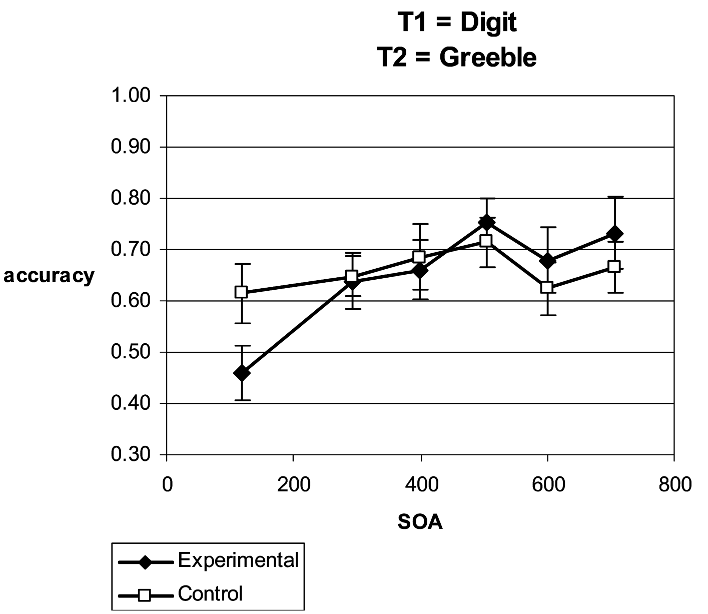

Fig. 12. Mean accuracy of second target (T2) report, given correct report of the first target (T1) as a function of stimulus onset asynchrony (SOA) and experimental condition in Experiment 7.

iment 7 stands in marked contrast to the long-lasting impairments that were observed when subjects discriminated the same T1 digits, but letter stimuli were in the T2 position.

The interference in Experiment 7 may have been short-lived because the observers made use of the configural information in the Greeble stimuli. These configural cues may have guided the discrimination of the Greeble stimuli even though the T1 digit task had occupied featural processing resources. However, this conclusion rests on two important assumptions. First, it must be shown that the masking in this procedure was adequate for overwriting the Greeble stimuli. Second, it must be verified that the discrimination of these Greebles was sufficiently demanding that we would be able to detect the effects of competition for a putative central resource. We tested both of these conclusions in Experiment 8 by placing faces in the T1 position, and Greeble stimuli in the T2 position. If Greeble discrimination in Experiment 7 was maintained because of access to a configural processing channel, then AB interference should be observed when faces are placed in the T1 position. Under these conditions, the T1 face should occupy all of the available channels for processing the second stimulus. However, if Greeble discrimination was preserved in Experiment 7 because of inadequate masking or because these stimuli do not require significant processing resources, then Experiment 8 should produce little evidence of AB interference.

11. Experiment 8

11.1. Methods

11.1.1. Observers

Eight volunteers from the local community were paid for their participation in a 1-h session. The observers were between the ages of 18 and 30 years, with normal or corrected-to-normal vision.

11.1.2. Procedure

Each observer participated in six blocks of 30 trials in the experimental condition and six blocks of 30 trials in the control condition. The order of these conditions was counterbalanced across observers. One block of practice was administered before each condition.

11.1.3. Stimulus presentation

All aspects of the procedure were identical to those of Experiment 6 with the following exceptions: Six SOAs were tested: 0, 235, 353, 470, 706, and 1176 ms. Greeble stimuli were placed in the T2 position. These stimuli were masked with the same masks that were used in Experiment 7. During a single practice block, the exposure duration for the Greebles was set to 141 ms. For the remainder of the experiment, the exposure duration for the Greebles was 94 ms.

11.2. Results and discussion

The mean accuracy for the first target was 80%. T2 accuracy is displayed in Fig. 13 as a function of condition and SOA. As Fig. 13 illustrates, there were long-lasting deficits in accuracy in the experimental condition relative to the control condition. A two-way ANOVA with subjects, conditions (experimental vs. control), and SOA as factors revealed a main effect of condition ($F[1,7] = 44.8, p < .01$), with lower accuracy in the experimental condition (54%) than in the control condition (73%). There was also main effect of SOA ($F[5,35] = 9.8$, $p < .01$), and a significant interaction of condition and SOA ($F[5,35] = 2.8, p < .05$); as SOA increased, there was a gradual convergence of accuracy in the experimental and control conditions. The substantial interference revealed in this experiment suggests that the paucity of AB interference in Experiment 7 cannot be attributed to inadequate masking of the Greebles or an easy Greeble discrimination task. Instead, we hypothesize that AB interference was observed in Experiment 8 because the face discrimination task occupied a configural processing channel that would have otherwise been available to guide the T2 Greeble discrimination. Of course, this interpretation also makes another clear prediction. T2 face stimuli should suffer from AB interference if Greebles are place in the T1 position. Experiment 9 tested this prediction. In addition, Experiment 9 provided an opportunity to demonstrate AB interference for faces even when the stimuli eliciting the interference were not also faces.

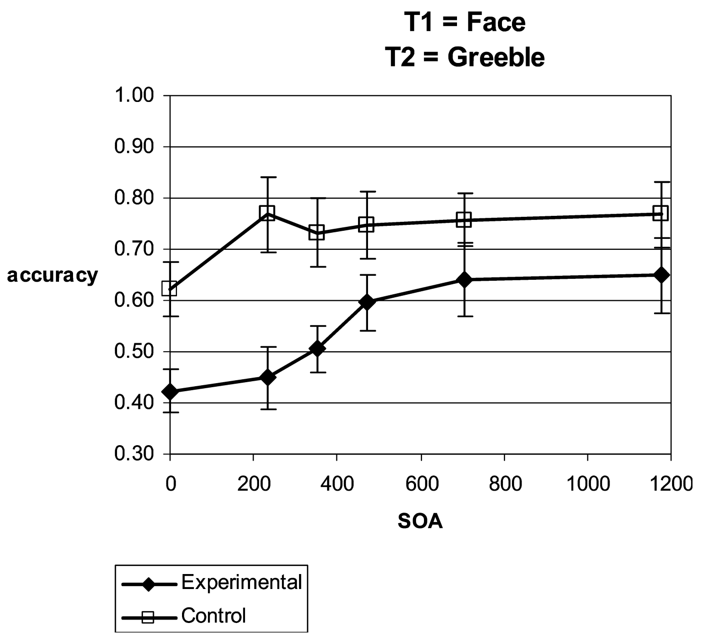

Fig. 13. Mean accuracy of second target (T2) report, given correct report of the first target (T1) as a function of stimulus onset asynchrony (SOA) and experimental condition in Experiment 8.

12. Experiment 9

12.1. Methods

12.1.1. Observers

Eight volunteers from the local community were paid for their participation in a 1-h session. The observers were between the ages of 18 and 30 years, with normal or corrected-to-normal vision.

12.1.2. Procedure

All aspects of this procedure were identical to those of Experiment 8, except that the order of the stimuli was reversed.

12.2. Results and discussion

The mean accuracy for the first target was 71%. T2 accuracy is displayed in Fig. 14 as a function of condition and SOA. As Fig. 14 illustrates, Greebles in the T1 position elicited long-lasting AB interference for the subsequent discrimination of faces. A two-way ANOVA with subjects, conditions (experimental vs. control), and SOA as factors revealed a main effect of condition ($F[1,7] = 14.1$, $p < .01$), with lower accuracy in the experimental condition (77%) than in the control condition (90%). There was also main effect of SOA ($F[5,35] = 6.4$, $p < .01$), and a significant interaction of

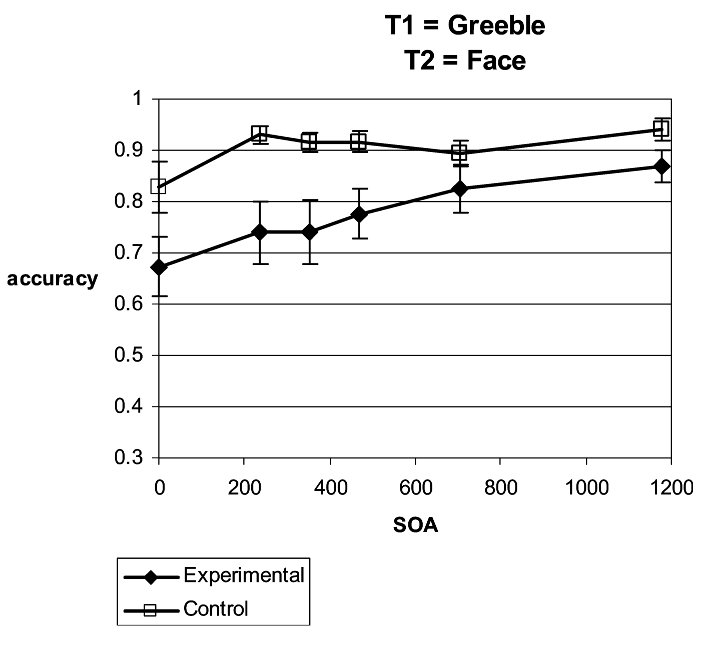

Fig. 14. Mean accuracy of second target (T2) report, given correct report of the first target (T1) as a function of stimulus onset asynchrony (SOA) and experimental condition in Experiment 9.

condition and SOA ($F[5, 35] = 2.5$, $p = .05$); as SOA increased, there was a gradual convergence of accuracy in the experimental and control conditions. AB interference for faces does not entail competition with other face stimuli. Instead, we suggest that interference can be predicted based on the processing overlap between the stimuli in the T1 and T2 positions.

13. General discussion

13.1. A multi-channel account of AB interference

The attentional blink has been demonstrated over an impressive range of task conditions, with virtually every class of stimulus showing reliable evidence of AB interference. Nevertheless, the current results suggest that the AB effect may not be the result of a central processing bottleneck in visual perception. We have observed that a T1 processing load (digit discrimination) that is capable of inducing long-lasting AB interference in letters may have no effect at all on the accuracy of face discrimination. These data are incompatible with the idea that AB interference reflects competition between objects for a single obligatory stage of visual processing. Instead, we hypothesize that multiple processing channels may be available for the discrimination of faces. The present results can be explained by assuming one channel that processes feature-based information, and another channel that processes configural information. This multi-channel account predicts that AB interference will be observed when T1 processing occupies every channel that is available for discriminating the second target. According to this hypothesis, face discrimination did not suffer from AB interference when T1 was a digit because the digits were discriminated using only the feature-based channel. Thus, the configural channel was still available for the accurate discrimination of faces during the AB period.

By contrast, a single-channel view of AB interference would have difficulty accounting for the fact that the T1 digit discrimination task induced long-lasting AB interference for letter discrimination, but not for face discrimination. These data therefore have implications for any model of AB interference that asserts competition between all visual stimuli for a single processing resource. For example, consider the two-stage model of AB interference proposed by Chun and Potter (1995). They suggest that each target passes through two distinct stages of processing. In the first stage of processing items are identified in a process that is relatively free of capacity limitations (with the presentation rates typical in the AB paradigm). However, without further processing, representations are left in a fragile state that is subject to rapid forgetting or overwriting by subsequent stimuli (Giesbrecht & Di Lollo, 1998). The second stage of processing is capacity-limited. During this postperceptual stage, the representations are consolidated into a durable form that that is capable of guiding subsequent behavioral responses. By this view, AB interference occurs because T2 cannot gain access to second stage processing. Several accounts of AB interference have been proposed that are broadly consistent with this two-stage model (e.g., Duncan et al., 1994; Jolicoeur, 1999). In each case, AB interference arises because of competition for a critical resource during a postperceptual stage of processing. The present results suggest that second stage processing may not be accomplished through a single obligatory resource for visual perception. Otherwise, any T1 processing load that causes AB interference for a subsequent visual stimulus (by occupying a mandatory resource for visual processing) should be capable of interfering with the identification of all classes of visual stimuli. Thus, the single resource view is challenged by the observation that face discrimination is unaffected by the same T1 task that causes long-lasting AB interference for letters.

This is not the first case in which AB interference was found to be reduced or eliminated. Shapiro et al. (1997) found that one's own name can be detected during the AB period without significant interference. However, this result is not necessarily incompatible with the single resource view, because personal names are significantly easier to process than other words. As Shapiro and colleagues noted, personal names may benefit from a lower threshold for activation, allowing them to be detected effectively with reduced resources. Therefore, even if personal names and other words compete for a common resource, one's own name might still be perceived during the AB period because it imposes a small processing load. By contrast, the present experiments provided four separate indications that the faces avoided AB interference while imposing a larger processing load than the letter discrimination task. First, in Experiments 1 and 2 the exposure duration of the letters and faces was identical (as was the T1 processing load), but accuracy in the control condition was lower for the faces than for the letters. Second, in Experiment 3 we matched the accuracy of face and letter discrimination (in the control condition) by setting the exposure duration of each stimulus category on a within-subject basis. This procedure replicated Experiments 1 and 2 and also revealed that the faces required significantly higher exposure durations in order to reach the same level of accuracy as the letters. Third, in Experiment 5 a more challenging set of masks (other faces) reduced face discrimination accuracy 27% relative to Experiment 2, but there was still no AB interference. Finally, in Experiment 4 we found that T1 faces induced a substantial attentional blink for digits, verifying that the face discrimination task places a significant load on AB-related processes. These observations converge to suggest that faces avoid AB interference because of qualitative rather than quantitative differences in processing load. More specifically, face discrimination may have been successful during the AB period because it had access to a configural processing channel that was not disrupted by the T1 digit task.

13.2. The scope of the present results

Of course, we do not mean to suggest that faces are somehow immune to dual task interference in general. Nor do these results suggest that face discrimination will always escape interference in attention-demanding tasks. Instead, these data can constrain models of the specific postperceptual process that is at the root of AB interference. There is clear evidence that the locus of AB interference is different, for example, from the perceptual stages of processing that are modulated by spatial attention. While spatial attention can affect the quality of the earliest stages of visual analysis, AB interference has its effects after a fully identified representation of the target stimulus has been formed. Vogel et al. (1998) demonstrated this point by measuring event-related potentials (ERPs) during the AB paradigm. They found that only a specific subset of target-evoked ERPs were modulated during the AB period. Early components of the ERP response that indexed perceptual stages of target identification were unaffected, while later components that indexed the updating of working memory were suppressed. A different profile of effects has been documented in ERP studies of spatial selection, in which there is clear evidence that attention can modulate perceptual processing (e.g., Mangun, Hillyard, & Luck, 1993). The contrast between the effects of spatial selection and AB interference underscores the important point that completely different cognitive processes may be at the core of interference in two tasks that reveal "attentional" limitations. Bearing this in mind, our conclusions are limited to models of the postperceptual interference that is observed during the AB period.

These considerations also motivated our use of the two-target paradigm instead of the more widely-used RSVP procedure. As discussed in the Introduction, the RSVP procedure may interfere with T2 processing in multiple ways. Challenges in the RSVP procedure include the need to inhibit distractor stimuli, the temporal uncertainty of T1 onset, and the use of independent stimulus features such as color to identify the target stimuli. Nevertheless, the AB effect has been the basis for several models that predict temporary impairments in T2 processing in the absence of these potential sources of interference. For example, the two-stage models discussed above predict that second stage resources will be occupied by the simple act of building a perceptual representation for T1 that is useful for overt responses. Thus, while there is clear evidence that distractor stimuli affect performance in the RSVP procedure (e.g., Isaak, Shapiro, & Martin, 1999), current models still predict that if the targets are adequately masked there should be AB interference in the absence of additional distractor objects. We argued in the Introduction that these models predict AB interference whenever three specific conditions are met. (1) The T1 task occupies sufficient processing resources to deny subsequent stimuli access to stage two processing. Experiments 1 and 3 verified this by showing that the T1 digit task induced long-lasting AB interference for subsequent letter targets. (2) The T2 task requires sufficient processing resources to show the effects of the T1 processing load. We have already discussed several indications that the faces were more difficult to discriminate than the letters, and that the face discrimination induces long lasting AB interference for digits and Greebles. (3) The first and second targets are adequately masked (e.g., Giesbrecht & Di Lollo, 1998; Grandison et al., 1997; Sieffert & Di Lollo, 1997). The masks we used in Experiments 2 and 3 were certainly sufficient to prevent visible persistence. In Experiment 5, we attempted to maximize the chances for object substitution by using another set of faces as masks for the face targets. This manipulation increased the difficulty of the task markedly, but no AB interference was observed. Finally, we observed AB interference with the same face photographs in Experiment 6 (by using a face as T1) and Experiment 9 (by using a Greeble as T1). These experiments show that the earlier observations of preserved face discrimination were not a product of insufficient masking procedures. Given that all three of these criteria were satisfied by our procedure, we conclude that the absence of AB interference for faces in Experiments 2, 3, and 5 is inconsistent with a single-channel view of AB interference. This forms the basis for our hypothesis that there are multiple channels available for stage two processing.

It is also useful to probe the implications of procedures that do find AB interference for faces. There are at least two such cases that would be consistent with the multi-channel account we have proposed for AB interference. First, there is the scenario that we have already tested in Experiments 6 and 9. If T1 processing occupies all of the channels that are available for face discrimination, then face discrimination should be impaired during the AB period. Thus, the multi-channel hypothesis would predict AB interference for faces with a range of T1 stimuli. Second, if the face targets are distinguished from each other only on the basis of features, then the configural processing channel would not support face discrimination. For example, using the same T1 digit task, pilot work in our lab has shown AB interference for faces that were distinguished only by mouth shape. This is consistent with the multi-channel hypothesis, because the only channel that could discriminate one target from another (i.e., the feature-based channel) was occupied by T1 processing. Finally, the contrast between the results of Experiments 2 and 4 showed that top-down settings can determine whether AB interference is observed or not. This raises the possibility that the constellation of processing channels occupied by a specific stimulus set are not determined in a wholly stimulus-driven manner. If 0, then AB interference might be observed with faces if observers were induced to adopt a feature-based strategy for discriminating these stimuli.

13.3. The role of top-down control

The observer's top-down settings play a key role in AB interference. This is apparent in the contrast between the results of Experiments 2 and 4. Even though the displays in these experiments were identical at the 0 ms SOA, both face and digit discrimination at this SOA was significantly better in Experiment 4. This may be a result of differences in the typical order of the stimuli in these experiments. The T1 digit appeared first in 90% of the trials in Experiment 2, while the T1 face appeared first in the 90% of the trials in Experiment 4. Thus, observers probably prepared to process digits in the Experiment 2 and faces in Experiment 4. Apparently, these two modes of preparation had a dramatic impact on the degree of dual task interference that was observed in identical visual displays. This aspect of the present results is also relevant to the question of whether all visual stimuli must access an obligatory resource for conscious perception. If the attentional settings of the observer can determine whether interference is observed or not, then the interference cannot be explained by competition for an obligatory, limited-capacity process.

14. Conclusions

The central theoretical issue in the present work—whether or not performance is limited by competition for a unitary resource—has also been extensively addressed in other dual task paradigms. For example, the psychological refractory period (PRP) paradigm, in which observers produce speeded responses to sequentially presented stimuli, has provided basic insights into the nature of dual task costs (for reviews, see Meyer & Kieras, 1997; Pashler, 1994). There has been vigorous debate, however, over whether such dual task costs are caused by the limited access to a central processing resource, such as response selection. Opposing this view, some have suggested that dual task interference is better explained by competition for multiple processing resources, each with a limited processing capacity (e.g., Navon & Gopher, 1979). By this account, dual task costs can be predicted based on the degree to which each task calls upon overlapping components of a broad range of resources.

In line with multiple resource theory, Wickens (1980) described a set of empirical patterns that are problematic for unitary resource accounts of dual task costs. We describe two of these patterns below, because they provide a useful framework for considering our own observations: (1) *Difficulty insensitivity* refers to an absence of reductions in task performance even when the difficulty of a concurrent task is increased. (2) *Structural-alteration effects* refer to cases in which the structural requirements of a primary task are changed—without reducing the difficulty of the primary task—and significant reductions in dual task interference are observed. Competition for a shared resource should produce interference whose magnitude is tightly linked to the difficulty of the competing tasks. Thus, changes in the structural requirements of one task without changes in its difficulty should not affect the level of dual task interference if both tasks compete for access to a single resource. When these empirical patterns are observed, one natural explanation is that the two tasks draw upon independent resources. The current experiments provide evidence of precisely this kind of disconnect between AB interference and the difficulty of the competing tasks. Experiments 1–3 showed that changing the T2 stimulus from a letter to a face resulted in significant improvements in both T1 and T2 accuracy. At the same time, there were multiple indications that the face discrimination was more difficult than the letter discrimination. Thus, the overall difficulty of T1 and T2 discrimination cannot account for the level of AB interference in these studies. Instead, we suggest that distinct constellations of processing structures are recruited by faces and alphanumeric stimuli, and that AB interference is predicted by the degree of overlap between these components.

The finding that a demanding face discrimination task can be performed without any impairment during the AB period suggests that AB interference may not be an example of a central processing bottleneck in visual perception. Instead, we suggest that there are multiple channels available for processing during the AB period. By this view, AB interference is contingent upon whether T1 occupies every processing channel that can accomplish the T2 discrimination. The single-channel view is consistent with observations of AB interference across an impressive range of stimuli. However, a single robust exception to the rule should be sufficient to cast doubt on this hypothesis. Future research directed towards understanding this exception should provide a better understanding of the boundary conditions of this limitation in sequential target processing.

References

Brehaut, J. C., Enns, J. T., & Di Lollo, V. (1999). Visual masking plays two roles in the attentional blink. *Perception & Psychophysics, 61*(7), 1436–1488.

Broadbent, D. E., & Broadbent, M. H. P. (1987). From detection to identification: Response to multiple target in rapid serial visual presentation. *Perception & Psychophysics, 42*, 105–113.

Chun, M. M., & Potter, M. C. (1995). A two-stage model for multiple target detection in rapid serial visual presentation. *Journal of Experimental Psychology: Human Perception and Performance, 21*, 109–127.

Di Lollo, V., Kawahara, J., Zuvic, S., & Visser, T. A. W. (2001). The preattentive emperor has no clothes: A dynamic redressing. *Journal of Experimental Psychology: General, 130*(3), 479–492.

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

Farah, M., Wilson, K., Drain, M., & Tanaka, J. (1998). What is "special" about face perception? *Psychological Review, 105*(3), 482–498.

Gauthier, I., & Tarr, M. J. (2002). Unraveling mechanisms for expert object recognition: Bridging brain activity and behavior. *Journal of Experimental Psychology: Human Perception and Performance, 28*(2), 431–446.

Gauthier, I., William, P., Tarr, M. J., & Tanaka, J. (1998). Training 'greeble' experts: A framework for studying expert object recognition processes. *Vision Research, 38*, 2401–2428.

Giesbrecht, B., & Di Lollo, V. (1998). Beyond the attentional blink: Visual masking by object substitution. *Journal of Experimental Psychology: Human Perception and Performance, 24*(5), 1454–1466.

Grandison, T. D., Ghirardelli, T. G., & Egeth, H. E. (1997). Beyond similarity: Masking of the target is sufficient to cause the attentional blink. *Perception & Psychophysics, 59*(2), 266–274.

Isaak, M. I., Shapiro, K. L., & Martin, J. (1999). The attentional blink reflects retrieval competition among multiple rapid serial visual presentation items: Tests of an interference model. *Journal of Experimental Psychology: Human Perception and Performance, 25*(6), 1774–1792.

Jolicoeur, P. (1999). Concurrent response-selection demands modulate the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance, 25*(4), 1097–1113.

Jolicoeur, P., & Dell' Acqua, R. (2000). Selective influence of second target exposure duration and task-sub-1 load effects in the attentional blink phenomenon. *Psychonomic Bulletin & Review, 7*(3), 472–479.

Joseph, J. S., Chun, M. M., & Nakayama, K. (1996). Attentional requirements in a 'preattentive' feature search task. *Nature, 379*, 805–807.

Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience, 17*, 4302–4311.

Luck, S. J., Vogel, E. K., & Shapiro, K. L. (1996). Word meanings can be accessed but not reported during the attentional blink. *Nature, 383*, 616–618.

Mangun, R. G., Hillyard, S. A., & Luck, S. J. (1993). Electrocortical substrates of visual selective attention. In D. Meyer, & S. Kornblum (Eds.), *Attention and performance XIV* (pp. 219–243). Cambridge, MA: MIT Press.

Meyer, D. E., & Kieras, D. E. (1997). A computational theory of executive cognitive processes and multiple-task performance: Part 1. Basic mechanisms. *Psychological Review, 104*, 3–65.

Navon, D., & Gopher, D. (1979). On the economy of the human-processing system. *Psychological Review, 86*, 214–255.

Pashler, H. (1994). Dual-task interference in simple tasks: Data and theory. *Psychological Bulletin, 116*, 220–244.

Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink. *Journal of Experimental Psychology: Human Perception and Performance, 18*, 849–860.

Ro, T., Russell, C., & Lavie, C. (2001). Changing faces: A detection advantage in the flicker paradigm. *Psychological Science, 12*(1), 94–99.

Ross, N. E., & Jolicoeur, P. (1999). Attentional blink for color. *Journal of Experimental Psychology: Human Perception and Performance, 25*(6), 1483–1494.

Shapiro, K. L., Arnell, K. M., & Drake, S. H. (1991). Stimulus complexity mediate4s target detection in visual attention search. *Investigative Ophthalmology and Visual Science, 32*, 1040.

Shapiro, K. L., Caldwell, J., & Sorensen, R. E. (1997). Personal names and the attentional blink: A visual "cocktail party" effect. *Journal of Experimental Psychology: Human Perception and Performance, 24*, 504–514.

Sieffert, A., & Di Lollo, V. (1997). Low-level masking in the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance, 10*(4), 1061–1073.

Tanaka, J. W., & Gauthier, I. (1997). Expertise in object and face recognition. In R. L. Goldstone, P. G. Schyns, & D. L. Medin (Eds.), *Psychology of learning and motivation* (36, pp. 83–85). San Diego, CA: Academic Press.

Tarr, M. J., & Gauthier, I. (2000). FFA: A flexible fusiform area for subordinate-level visual processing automatized by expertise. *Nature Neuroscience, 3*(8), 764–769.

Vogel, E. K., Luck, S. J., & Shapiro, K. L. (1998). Electrophysiological evidence for a postperceptual locus of suppression during the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance, 24*, 1656–1674.

Ward, R., Duncan, J., & Shapiro, K. (1997). Effects of similarity, difficulty, and nontarget presentation on the time course of visual attention. *Perception & Psychophysics, 59*, 593–600.

Wickens, C. D. (1980). The structure of processing resources. In R. Nickerson, & R. Pew (Eds.), *Attention and performance VIII* (pp. 239–257). Hillsdale, NJ: Erlbaum.