A Two-Stage Model for Multiple Target Detection in Rapid Serial Visual Presentation

Marvin M. Chun and Mary C. Potter
Massachusetts Institute of Technology

When 2 targets are presented among distractors in rapid serial visual presentation, correct identification of the 1st target results in a deficit for a 2nd target appearing within 200–500 ms. This attentional blink (AB; J. E. Raymond, K. L. Shapiro, & K. M. Arnell, 1992) was examined for categorically defined targets (letters among nonletters) in 7 experiments. AB was obtained for the 2nd letter target among digit distractors (Experiment 1) and also for a 3rd target (Experiment 2). Results of Experiments 3–5 confirmed that AB is triggered by local interference from immediate posttarget stimulation (Raymond et al., 1992) and showed that AB is modulated by the discriminability between the 1st target and the immediately following distractor. Experiments 5–7 further examined the effects of both local interference and global discriminability. A 2-stage model is proposed to account for the AB results.

Researchers working on visual attention have focused on capacity limitations that arise when multiple stimuli must be processed in a single spatial array. Different issues arise when stimuli are presented sequentially. In this study, we examined attentional limitations for processing a temporal sequence of visual stimuli. When participants search for targets among stimuli presented in a sequence at high rates, correct identification of one target produces a marked deficit for detecting a subsequent target appearing in a 200–500 ms interval after the onset of the first one (Broadbent & Broadbent, 1987; Raymond, Shapiro, & Arnell, 1992). These tasks involve the use of rapid serial visual presentation (RSVP), in which each item replaces the previous one at the same spatial location. The RSVP paradigm has been a useful tool for researchers exploring the temporal characteristics of information processing because it provides the experimenter with precise control not only over the time a given item is in view, but also over the preceding and subsequent processing demands on the participants. In RSVP each item not only eliminates the previous item from sensory storage (Kahneman, 1968), but also presents a new item to be processed, thus constraining the time available for higher level cognitive as well as perceptual processing (Potter, 1976).

One of the earliest studies that used RSVP to study target detection was that of Lawrence (1971), who measured recognition accuracy for a single capitalized target word appearing at various serial positions within an RSVP stream of lowercase distractor words presented at rates between 6 to 20 items per second. Recognition accuracy decreased with higher rates of presentation. Target identification errors mainly consisted of distractor word intrusions (78%); of these intrusions, 88% were of the word directly following the target. The prevalence of posttarget intrusion errors suggests that a two-step filtering process exists in which the participant notes the target-defining capitalization and then encodes a word (Gathercole & Broadbent, 1984; Lawrence, 1971; McLean, Broadbent, & Broadbent, 1983).

The Attentional Blink

The focus of this article is on interference with subsequent target detection after a first target has been identified. Broadbent and Broadbent (1987) extended Lawrence's (1971) study by requiring participants to report two uppercase words embedded among lowercase words, with the lag between the two targets varied. Using presentation times of 80 to 120 ms per item, which allowed a single target to be identified on a high proportion of the trials, Broadbent and Broadbent showed that correct identification of the first target (T1) interfered with the identification of a second target (T2) appearing within 500 ms. They demonstrated a similar pattern of deficits in a condition where the targets could be a word appearing in uppercase (defined by a physical feature) or an animal word (categorically defined target). They concluded that detection of the features that signal the presence of a target is possible at the rates they used, but once a target is detected, it triggers more demanding processes of identification that interfere with subsequent target detection and identification. They cited the work of

Duncan (1980), who pointed out that targets in simultaneous arrays interfere with each other.

In a different task, Weichselgartner and Sperling (1987) presented a stream of digits in RSVP, and participants were asked to report the four digits starting with a target digit that was cued by luminance or the outline of a square. The data were collapsed across trials to produce a distribution of recall probability as a function of lag from the target. A bimodal distribution was obtained, with report most likely for digits appearing either 0–200 ms or 400–500 ms after the onset of the target cue. Thus, there was an interval between 200 and 400 ms after the initial target in which participants seemed unable to encode further targets. The authors interpreted this bimodal distribution as indicative of two different modes of attention: a quick, effortless, automatic process triggered by target detection and a slower, effortful, controlled process, the latency of which depended on practice and task difficulty.

Raymond, Shapiro, and Arnell (1992) replicated and extended these results using a different procedure in which all of the items were single letters. T1 was marked by being white, whereas the other letters were black against a gray background. T2 (presented on 50% of the trials) was always the letter X; the task was to identify T1 and to decide whether the X probe had occurred or not. When the X appeared with a stimulus onset asynchrony (SOA) between 200 and 500 ms, it was often missed. Raymond et al. hypothesized that in their experiments the processing time needed for identification of a target item exceeded the onset-to-onset time of 90 ms. They proposed that the onset of a new stimulus before the processing of T1 is complete causes interference and invokes an attentional suppression mechanism. Attentional suppression, characterized as a shut-and-locked attentional gate, lasts for several hundred milliseconds, increasing the probability that targets appearing within that interval will be missed. They termed this deficit an *attentional blink* (AB), in analogy to the suppression of visual processing during rapid saccadic eye movements (Volkman, Riggs, & Moore, 1980). Raymond et al. tested an implication of their hypothesis, that adding processing time immediately after the first target would eliminate AB. As they predicted, when a blank interval of 90–270 ms was substituted for the items that immediately followed T1, no AB deficit was found.

Like Weichselgartner and Sperling (1987), Raymond et al. (1992) found that when the next item after T1 was T2 (at Lag 1), then T2 was relatively easy to detect. Thus the AB curve was again a U-shaped function of lag. Raymond et al. proposed that the deficit in the 200–500 ms SOA range is the result of the shutting of an attentional gate when the item following T1 interferes with its identification. Accordingly, the gate is closed to prevent further confusion. Thus, T1 and an immediately following item are processed together, so that a second target at Lag 1 is not suppressed.

Thus, two previous explanations of the AB deficit are that it results from an attentional suppression mechanism or that it reflects two distinct types of attentional processes. We propose a new two-stage model that incorporates aspects of both of the previous models. In this two-stage model, the AB deficit is the result of a limited-capacity second stage in which targets detected in the first stage are processed and consolidated serially. This second processing stage does not begin until the first stage signals the probable presence of a target, so that the second stage typically overlaps with the presentation of the following items in an RSVP sequence. This stage of processing operates on whatever internal representations of items are available at the time processing begins, and thus not only the first target but also the immediately following item (target or distractor) are likely to be included in processing. When the duration of this limited-capacity processing of T1 and the following item exceeds the SOA between T1 and T2, interference with the processing of T2 results. When the Lag 1 item (the item that immediately follows T1) is a distractor, the difficulty of discriminating it from T1 determines the time course of the second-stage bottleneck. When the Lag 1 item is also a target (T2), then it is processed together with T1 and is likely to be reported.

In this study, we looked more closely at AB and the conditions necessary for its appearance, focusing on the nature of the nontarget items and in particular the item (if any) that immediately follows T1. In all experiments the task was to detect and identify unspecified letters among digits or symbols; the lag between the letter targets varied. In the general discussion, we consider how the AB deficit in RSVP can be understood in terms of the two-stage model outlined above, with a first stage of target detection and a second capacity-limited stage in which candidate targets are fully identified and registered in memory. We also discuss how the AB deficit relates to previous research on dual-target and dual-task interference, and what important differences arise from the characteristics of the RSVP paradigm.

Experiment 1

In the first experiment, we replicated and extended previous studies of the AB deficit (Broadbent & Broadbent, 1987; Raymond et al., 1992; Weichselgartner & Sperling, 1987) using a somewhat different task. In most earlier studies, targets in an RSVP stream were specified by a separate target-defining feature such as color, lettercase, or luminance.[1] The target-defining features of T1 were independent of the target features to be reported, requiring a conjoining of the two sets of features. Such conjoining of arbitrary features is believed to require focal attention (Treisman & Gelade, 1980), and it is possible that AB results from such focal attention to T1. However, AB may be the result of attentional processing inherent to identification and consolidation of T1 appearing in RSVP.

---

[1] In one of their experiments, Broadbent and Broadbent (1987) used targets that were defined by their categorical identity (animals among nonanimals). However, these categorically defined targets were intermixed with targets defined by a separate physical feature (uppercase words among lowercase distractors), making it difficult to draw conclusions about pairs of categorically defined targets.

To evaluate these alternative hypotheses, we constructed the present experiment so that the targets to be detected were defined by their categorical identity as letters. Duncan (1980, 1983) proposed that the categorical identities of letters and digits are available preattentively and may serve as a basis for selection into a more limited-capacity system. Others have suggested that digits and letters may be separated on the basis of some key stimulus features (Krueger, 1984; Treisman & Gelade, 1980). Still others (Sperling, Budiansky, Spivak, & Johnson, 1971) have suggested that the specific identity of digits (and presumably also letters) is available at least as quickly as their categorical identity. Whether the full identity of an item is available preattentively or not, the features used to signal targethood in the present experiments were the same features that needed to be analyzed for subsequent report.

Another characteristic of many of the previous dual-target experiments was that the target specification for T1 was different (or might have been different) from that for T2. It is possible, therefore, that the needed switch in set from T1 to T2 would cause AB. In this study, however, participants were simply instructed to report the two letters appearing in the stream of digits. Thus, unlike in Raymond et al.'s (1992) probe detection task, participants did not need to change selection set from T1 to T2 (e.g., search through black nontargets for a white letter [T1] and then a black letter X [T2]). Therefore, the first question was whether an AB deficit for T2 would appear in a categorically defined target task that would not require conjunction of arbitrary features and would not require the participant to switch set from T1 to T2.

Method

Participants. Six participants were involved in Experiment 1. In this and the later experiments, the participants were from the Massachusetts Institute of Technology (MIT) volunteer participant pool. All of the observers reported normal or corrected-to-normal visual acuity. Informed consent was obtained before participation, and observers were paid for their participation. None of the participants was aware of the purpose of these experiments.

Design and procedure. The stimuli were 8 single digits and 24 capital letters (0, 1, O, and I were omitted to avoid confusion). Each trial consisted of 16 items (14 digits and 2 letters). The digits on a given trial were randomly generated by the computer under the constraint that the same digit could not appear in the previous four positions. Two randomly sampled (but not identical) uppercase letters, designated T1 and T2, were selected as targets. The position of T1 was randomly permuted so that it appeared an equal number of times in Serial Positions 3–7. Eight lags between T1 and T2, from Lag 1 (no intervening items, SOA = 100 ms) to Lag 8 (SOA = 800 ms), were crossed with the five serial positions of T1, and the design was replicated six times for a total of 240 trials, with 30 trials at each lag. One practice block of 20 trials was followed by three experimental blocks of 80 trials each.

The experiment was self-paced. The participant began each trial by pressing the space bar on the computer keyboard. A plus sign lasting 400 ms appeared at the center of the monitor screen for fixation. One hundred ms after the fixation cross went off, the stream of stimuli appeared successively without interstimulus blanks at the same location for 100 ms each (presentation rate =

10 items per second). The sequence was followed by an ampersand mask for 100 ms, signaling the end of the trial. Participants were instructed to report the two letters aloud immediately after the trial. They were encouraged to avoid making wild guesses. The experimenter wrote down the response; no feedback was given. The experiment was carried out in normal room illumination held constant for all participants.

Apparatus. The same experimental apparatus was used for all of the experiments presented in this article. The letter and digit stimuli were generated by an IBM-AT computer on a CRT screen with a rapid fade phosphor. The stimuli measured about 0.3 cm in width and 0.4 cm in height. The display was viewed from 30 cm; thus, each stimulus subtended about 0.57 × 0.76 degrees of visual angle. Figure 1 shows the stimulus sets.

Results and Discussion

In brief, a marked AB deficit, comparable with that reported by Raymond et al. (1992), was obtained in report of the second of two letter targets among digits. Thus, neither an arbitrary T1 specifier that must be conjoined with other target features nor a switch in target set from T1 to T2, is required to produce AB.

Columns 3–8 in Table 1 contain, for each lag, the percentage of trials on which neither T1 nor T2, only T1, only T2, and both T1 and T2 were correctly reported and on which T1 was missed. Letters were counted as correct regardless of the order of report. The next column contains the conditional percentage of T2 report given that T1 was not reported (T2|$\overline{T1}$), and the final column contains the conditional percentage of recall of T2 at each lag—that is, the percentage of trials on which T2 was reported given that T1 was reported (T2|T1). This conditional percentage was computed for each participant at each lag and was averaged across participants. This conditional percentage was used to measure the AB deficit, shown graphically in Figure 2. Such figures are used to present the main results of all of the experiments in this article.

Correct identification of T1 produced a lag-dependent deficit for reporting T2. A one-way analysis of variance

LETTERS : A B C D E F G H

J K L M N P Q R

S T U V W X Y Z

DIGITS : 2 3 4 5 6 7 8 9

SYMBOLS : < > = # % ? / *

Figure 1. Stimuli used in the present experiments as they appeared on the CRT screen.

Table 1
Full Summary of Target Report Percentages for Experiment 1

Lag	SOA (ms)	Neither T1 nor T2	Only T1	Only T2	T1 missed	Both T1 and T2	T2|T̄1	T2|T1
1	100	3	16	31	33	51	94	76
2	200	7	53	12	19	28	63	34
3	300	6	60	5	11	30	45	32
4	400	7	40	16	22	38	73	47
5	500	7	23	12	19	58	63	70
6	600	6	18	6	12	70	50	79
7	700	2	20	10	12	69	83	77
8	800	3	16	7	10	74	70	81

Note. SOA = stimulus onset asynchrony; T1 = first target; T2 = second target.

(ANOVA) on the conditional report of T2|T1 showed a significant effect of lag, $F(7, 35) = 13.5, p < .001$. Report of T2|T1 was lowest at Lags 2 and 3, improving with increasing lag. The blink function reached an asymptote at Lag 6. We refer to this U-shaped curve, especially the monotonic portion showing increasing performance with increasing SOA, as the *lag effect*.

As noted earlier, an important characteristic of the dual-target paradigm is that the AB posttarget deficit does not occur for a target immediately following T1. As Figure 2 shows, when T1 and T2 were temporally adjacent (Lag 1), T2 was reported on a high proportion (76%) of the trials on which T1 was correctly identified. The conditional probability of T2|T1 at Lag 1 was significantly higher than at Lag 2, $t(5) = 6.6, p < .01$, Lag 3, $t(5) = 4.1, p < .01$, and Lag 4, $t(5) = 2.1, p < .05$, one-tailed. (All p values for the t tests are two-tailed unless otherwise specified.)

However, the results in Table 1 show that T1 was missed on a higher proportion of trials (33%) when it was immediately followed by a second target than when T2 appeared at any other lag (for all t-test comparisons, $p < .05$, one-tailed). That is, T1 was less likely to be missed when it was followed by a distractor digit than when followed by another target letter. Thus, target identification in RSVP tasks appears to be affected by the kind of item that follows the

target. We explore the role of the item directly following the target in more detail throughout this article.

The results of Experiment 1 closely mirror those of earlier studies (Broadbent & Broadbent, 1987; Raymond et al., 1992; Weichselgartner & Sperling, 1987), showing a deficit for reporting a second target appearing in the interval 200–500 ms from onset of the first target when T1 has been identified and recalled correctly, while showing no deficit for an item appearing at Lag 1. The results demonstrate that the AB deficit occurs for targets that are categorically defined. Thus, AB is not merely a consequence of having to conjoin an arbitrary, separate, target-defining feature (such as color or lettercase) with target-identifying features (the particular target letter or word that is to be reported). Processing the features intrinsic to the identity of an item (here, a letter) and encoding the target for subsequent report were sufficient to produce interference in detecting a subsequent target.

The presence of a characteristic AB deficit in our task also supports the conclusion that it is not a switch in set from T1 to T2 that causes the deficit. Horlitz, Johnston, and Remington (1992) reported results that support this conclusion. They independently varied T1 and T2 to be either a color-specified letter identification task or a visual discrimination task (determine the orientation of a Landolt C). They showed that the AB deficit is caused neither by filter-switching costs (AB occurred whether the target-specifying color was the same or different for T1 and T2) nor by task-switching costs (AB occurred whether the tasks for T1 and T2 were the same or different). Raymond et al. (1992) also suggested that the absence of AB for Lag 1 in their task indicates that a switch in set is not likely to be a factor.

The results of Experiment 1 suggest that the categorical identity of each item appearing in RSVP is rapidly available (within 100 ms, the presentation rate used here). Apparently, detection of a target in the designated category triggers further processing of the item, and this subsequent stage of processing results in a temporary deficit for subsequent targets.
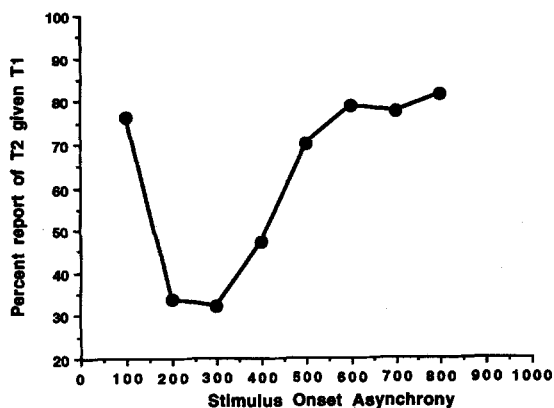

Figure 2. The attentional blink effect: Percentage report of the second target (T2; given report of the first target; T1) as a function of stimulus onset asynchrony in Experiment 1.

Experiment 2

In Experiment 2, we asked whether an AB deficit is produced not only after the first target (T1) but also after the

second target (T2) if the second target is processed success-fully. If AB is the result of a limited-capacity second stage for processing targets, as we propose, then it should occur *with* each detected target. Previous studies have not addressed this question directly. Weichselgartner and Sperling (1987) hypothesized that the onset of their second stage, a sustained, controlled mode of attention, allows for multiple subsequent items to be processed. They asked participants to report the four items beginning with detection of a target. Unfortunately, their presentation of the data does not reveal the contingent pattern of item report within a single trial. It is that pattern which needs to be examined to determine whether registration of each successive item is accompanied by an AB or whether their hypothesized sustained component, once initiated, allows for interference-free processing of several to-be-reported items. Regardless of the pattern of results in Weichselgartner and Sperling's procedure, report of an intermittent series of targets might be expected to generate repeated AB.

In Experiment 2, we instructed the participants to detect and report three letters embedded in an RSVP stream of digits. We systematically varied the lag between T1 and T2 and between T2 and T3.
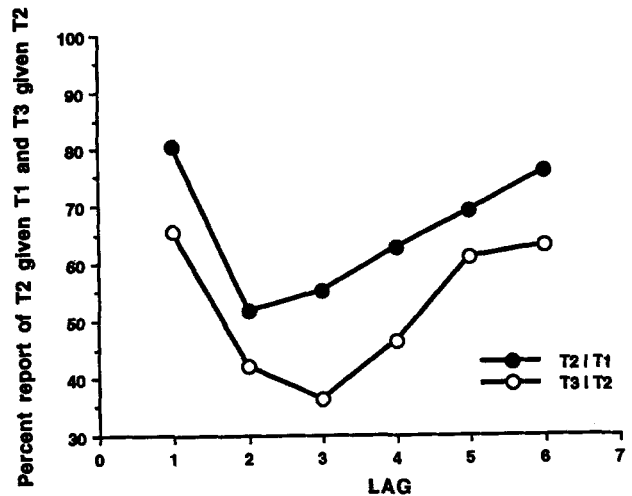

Figure 3. Attentional blink for three targets in Experiment 2: Percentage correct for the second target|the first target (T2|T1) and the third target|the second target (T3|T2) shown as a function of lag between T1 and T2 and between T2 and T3, respectively.

Method

The method used was identical to that of the previous experiments, except as specified below.

Participants. Sixteen participants were tested in this experiment.

Design and procedure. Three target letters were embedded in an RSVP stream of 17 digits. The position of T1 was randomized to appear in either Serial Position 2, 3, or 4. There were 6 possible lags (SOA = 100–600 ms) between T1 and T2 and between T2 and T3, making a total of 36 lag configurations. There were 8 trials in each of these 36 configurations, for a total of 288 trials. The experimental trials were preceded by a practice block of 20 trials. The entire experiment took about an hour and was self-paced, with breaks initiated by the computer after every 40 trials. The participants wrote their responses on a separate answer sheet.

Results and Discussion

For this and subsequent experiments, we omit the full table of results like that given for Experiment 1 and mainly focus on the conditional report of T2 and T3. Other results are reported where relevant. (Tables are available from Marvin M. Chun.) The main results are presented in Figure 3, in which the lag is indicated between T1 and T2 and between T2 and T3 for report of T2|T1 and T3|T2, respectively. The conditional probability of reporting T2 given that T1 was reported showed the standard AB deficit as a function of lag. The conditional probability of reporting T3 given that T2 was reported (regardless of T1–T2 lag) also showed a strong AB pattern. There was a main effect of lag for T2|T1, $F(5, 75) = 21.4$, $p < .001$, and T3|T2, $F(5, 75) = 18.4$, $p < .001$. Performance for reporting T2|T1 was better than for reporting T3|T2, $F(1, 15) = 83.3$, $p < .001$, but this difference did not interact with lag ($p > .22$).

Because of the small number of trials on which T1 was missed (16%), it was not possible for us to test the effect of T1 report on T3 across individual participants, but inspection of the two grouped curves for T3|T2 and T3|(T1&T2) indicated that the curves overlapped almost exactly, suggesting that T1's effect on T3 was minimal. The family of separate T3|(T1&T2) curves for each T1–T2 lag all revealed a U-shaped AB pattern, arguing against a sustained component of attention in our task. The group data analysis also showed that the probability of reporting T3 was lower when T2 was reported than when T2 was missed. All of these results support the main conclusion that each target occurrence creates an AB deficit for a subsequent target.

This demonstration of a second characteristic AB effect after a correctly identified T2 makes it clear that AB occurs after each selected target, not only after the first one. The decrement in performance for T3|T2 compared with T2|T1 may be due to memory load or to other factors, but the characteristic U shape of the deficit for T3|T2 and the lack of an interaction with T2|T1 suggest that both curves represent a common AB effect. This finding of a second AB effect does not directly disconfirm Weichselgartner and Sperling's (1987) proposal for a sustained mode of attention in their task, which required report of a continuous sequence of items once T1 was detected. Whereas our task could be described as involving partial report, Weichselgartner and Sperling's task used whole report for subsequent items. There may be differences between the two strategies for reporting items in RSVP. The results of Experiment 2 suggest that, at least in RSVP search tasks with targets embedded among hard-to-discriminate distractors, the detection and processing of each target creates a transient deficit for a subsequent one. This result supports the proposed two-stage model, inasmuch as each successive target creates a temporary second-stage bottleneck.

Experiment 3

In Experiments 1 and 2, the presence of AB in our categorically defined target task suggested that AB is not due to the conjoining of a target-defining feature with a target's identity, nor is it due to task switching between T1 and T2. In the next two experiments, we examined Raymond et al.'s (1992) proposal that the posttarget processing deficit is triggered by local events that interfere with target identification. The authors demonstrated that the AB deficit was reduced or eliminated when the first target (plus its normal interstimulus interval [ISI]) was immediately followed by a blank interval of 90 ms or longer, replacing one or several subsequent items (see Raymond et al., 1992, Experiments 3 and 4). In contrast, a blank interval appearing after a single immediate posttarget event did not prevent AB. Thus, an immediate posttarget event is what seemed to trigger the deficit for T2. Although Raymond et al. showed a main effect of blank duration, inspection of their results suggests that a blank of 90 ms (the duration of one item) was sufficient to significantly reduce the deficit for T2. In Experiment 3, we examined the effect of a single blank interval on report of letter targets among digits. If there is a reduction of the AB deficit for T2 when T1 is followed by a blank in our task, that would support the idea that the item immediately following T1 interferes with its identification and consolidation, thus increasing the duration of second-stage processing.

Method

Participants. Eight participants were tested in this experiment. One participant was replaced because of an error rate that exceeded a preset criterion (no more than 10% of trials on which neither target was reported correctly).

Design and procedure. Except as noted below, the method used was identical to that of Experiment 1. Each trial consisted of 15 items. T1 appeared randomly in Serial Positions 3–7, and T2 appeared at Lags 1–7 (Lag 8 was eliminated). There were two blank conditions and a baseline condition. In the baseline condition (equivalent to Experiment 1), no blanks were presented. The Lag 1 blank condition had a blank at Lag 1, with T2 appearing equally often at each of the other lags. The Lag 2 blank condition had a blank at Lag 2; T2 appeared equally often at all other lags. Thus, there were seven T2 lags in the baseline condition and six lags in each of the two blank conditions. There were 15 trials at each lag in each condition, for a total of 285 trials. The experiment was preceded by a block of 20 practice trials. Trials from all three conditions were intermixed within blocks and within participants.

Results and Discussion

Figure 4 shows the percentage report of T2 given correct report of T1. The baseline condition without a blank resulted in an AB pattern very similar to that found in Experiment 1. AB was reduced when a blank was presented immediately after T1, but not when the blank appeared at Lag 2. The two blank conditions were compared with the baseline condition in separate ANOVA tests. The lags included in comparison with the baseline condition were those
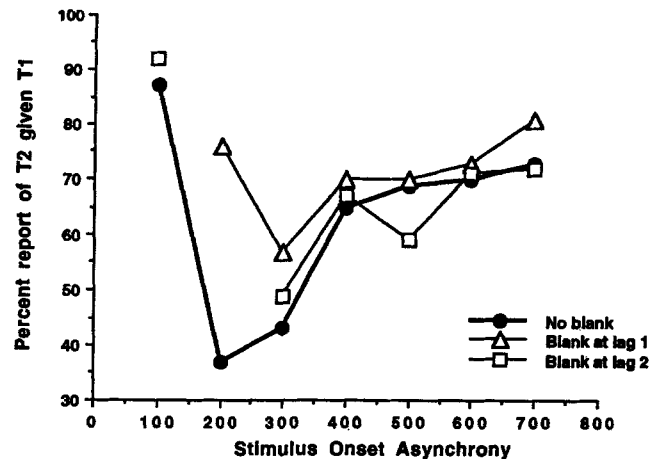

Figure 4. Effect of a blank among digits: Experiment 3.

lags used in the respective blank conditions. For the Lag 1 blank and the baseline conditions, there were significant main effects of condition, $F(1, 7) = 29.3, p < .001$, and lag, $F(5, 35) = 7.8, p < .001$. The interaction between condition and lag was also significant, $F(5, 35) = 3.9, p < .01$. Individual t tests at each lag showed that the Lag 1 blank condition led to better T2 performance than did the no-blank baseline at Lag 2, $t(7) = 5.8, p < .001$, and at Lag 3, $t(7) = 3.4, p < .01$, but not at Lags 4–7. For the Lag 2 blank condition versus the baseline, the ANOVA showed no difference between the two conditions ($p > .60$). There was a significant main effect of lag, $F(5, 35) = 10.2, p < .001$, reflecting the usual AB pattern, but the interaction between lag and condition was not significant ($p > .30$).

T1 was correctly reported on 87% of trials in the baseline condition, excluding the Lag 1 condition. Report of T1 was significantly better (98%) in the Lag 1 blank condition, $F(1, 7) = 20.0, p < .01$, but significantly worse (82%) in the Lag 2 blank condition, $F(1, 7) = 9.8, p < .05$. When T1 was followed by T2 at Lag 1, T1 was reported on 76% of the trials in the baseline condition and on 68% of the trials in the Lag 2 blank condition (in which T2 was followed by a blank). Thus, performance on T1 was higher when followed by a digit than when followed by another target (as in Experiment 1) and was best when followed by a blank.

The results are generally consistent with the findings reported by Raymond et al. (1992). A blank equivalent to the duration of an item (100 ms) was sufficient to reduce the deficit for T2 when the blank immediately followed T1, but not when it followed the first distractor after T1. Thus, the deficit for identifying T2 occurred only, or most strongly, when immediate posttarget stimulation interfered with T1 processing.

However, the release from AB in the present experiment does not appear to be complete. There was a suggestion of a deficit for targets presented after the blank interval that immediately followed T1 (as shown by the triangles in Figure 4). There was also a significant dip from an SOA of 200 ms, right after the blank, to an SOA of 300 ms after one

intervening digit, $t(7) = 3.6$, $p < .01$, as though the intervening digit produced a small-scale AB effect. That is, there was a relative benefit for a target at Lag 2, right after the Lag 1 blank interval, compared with a target at Lag 3. Raymond et al.'s (1992) results seem to show the opposite pattern, with an apparent deficit for items appearing immediately after the blank interval (the significance of this slight deficit is not reported). These minor discrepancies in results may be due to several differences in the methods of the two experiments. In Raymond et al.'s experiments, the RSVP stimuli were presented in flicker fashion, with blank intervals appearing between every item; in the present experiment, each new item replaced the previous one, with no ISI. Also, the basic tasks in the two experiments were different. In addition, we used only one blank duration (100 ms) in Experiment 2, whereas their experimental blocks consisted of blank durations of variable length (from 90 ms to 270 ms). Although our blank of 100 ms was sufficient to produce a significant benefit in report of T2, perhaps the duration was not long enough to allow Stage 2 processing of T1 to be fully completed before poststimulation resumed at Lag 2, resulting in the minor AB pattern just noted.

Any proposal for explaining the AB deficit needs to account for the fact that performance in detecting T2 is modulated by the sequence of visual events following the correctly identified T1. Among the posttarget events, the presence of an item directly following T1 appears to be most critical. A question not answered by the present experiment and by that of Raymond et al. (1992) is what characteristics of the visual event following T1 determine the presence and extent of the AB deficit. In Raymond et al.'s study, the interfering item was a letter that could be confused with the T1 letter, even though T1 appeared in a different color (white rather than black). Indeed, in most previous studies the item appearing immediately after T1 was always a member of the set of possible targets and thus could be confused with T1. However, in our experiments the interfering item was a digit rather than a letter, and even though the interfering item was not a potential target, it still produced a large deficit for reporting T2.

These findings raise the question of whether the source of interference can be any visual event or only a visual event that shares confusing visual features or categorical similarity with T1. Raymond et al. (1992) suggested in their general discussion that the amount of attentional suppression should be a function of the degree of visual similarity between T1 and the following item. In Experiment 4, we directly examined the effect of local interference by manipulating the ease with which an observer could discriminate between T1 and its following distractor and between T2 and its following distractor.

Experiment 4

In Experiment 4, we used distractor sets of digits and keyboard symbols to examine the effect of local interference. The keyboard symbols were chosen to be more discriminable from the letter targets than the digit distractors.

The letter targets were presented among a distractor set consisting of a mixture of digits and symbols, and we varied the type of distractor that immediately followed T1 and the distractor that immediately followed T2. Thus, the global discriminability of targets from distractors was constant across trials, but the items immediately following each target were varied.

Method

The method was identical to that of Experiment 1, except as noted below.

Participants. Sixteen participants were tested in this experiment.

Design and procedure. Only uppercase letters were used as targets; the letters I, O, U, V, and L were excluded. Digit characters two through nine (8 items) were used for the digit distractor set, and eight different keyboard symbols ($<$, $>$, $=$, $\#$, $\%$, $?$, $/$, and $*$) were used as distractors for the symbol distractor set. Figure 1 shows the stimuli. Trials included 13 items: T1, T2, and 11 distractors. T1 was randomly selected to appear in Serial Position 3, 4, or 5. T2 appeared at one of six SOAs (100–600 ms). The final item before the ampersand mask was always a distractor. The type of distractor (digit or symbol) immediately following T1 (which we denoted as D1) and the distractor type immediately following T2 (denoted as D2) were systematically varied. All other distractors were chosen randomly, with an equal probability of being either a symbol or a digit. There were 12 trials for each of the four possible D1 and D2 distractor combinations at each lag, for a total of 288 trials. D1 was not varied for Lag 1 (SOA $= 100$ ms) since T1 was directly followed by T2. The results for Lag 1 were analyzed separately. The experiment was preceded by a practice block of 20 trials.

Results and Discussion

The results for Experiment 4 by condition are shown in Figure 5. The effects of D1 and D2 are shown separately in parts A and B of Figure 6, respectively. Note that at Lag 1 there was no D1, so only D2 was a variable. Excluding Lag 1, there was a main effect of lag, $F(4, 60) = 15.5$, $p < .001$, type of D1, $F(1, 15) = 30.5$, $p < .001$, and type of D2, $F(1, 15) = 43.8$, $p < .001$. Neither interaction of D1 with lag nor D2 with lag was significant ($p > .21$). Thus, the results indicate that T2|T1 performance is sensitive to both the type of distractor (D1) immediately following T1 and the type of distractor (D2) following T2, with higher performance occurring when either target is followed by a symbol rather than a digit. The effects of D1 and D2 appear to be independent of each other, as suggested by the lack of interactions, either between D1 and D2 ($p > .82$) or among D1, D2, and lag ($p > .83$). For Lag 1, T2|T1 was reported on 89% of the trials when T2 was followed by a symbol and on 82% of the trials when T2 was followed by a digit, $F(1, 15) = 4.9$, $p < .05$.

The type of D1 had an effect on the proportion of trials on which T1 was correctly identified (excluding Lag 1). When T1 was followed by a symbol, it was reported on 90% of the trials, and when T1 was followed by a digit, it was reported on 79% of the trials, $F(1, 15) = 30.1$, $p < .001$. At Lag 1,
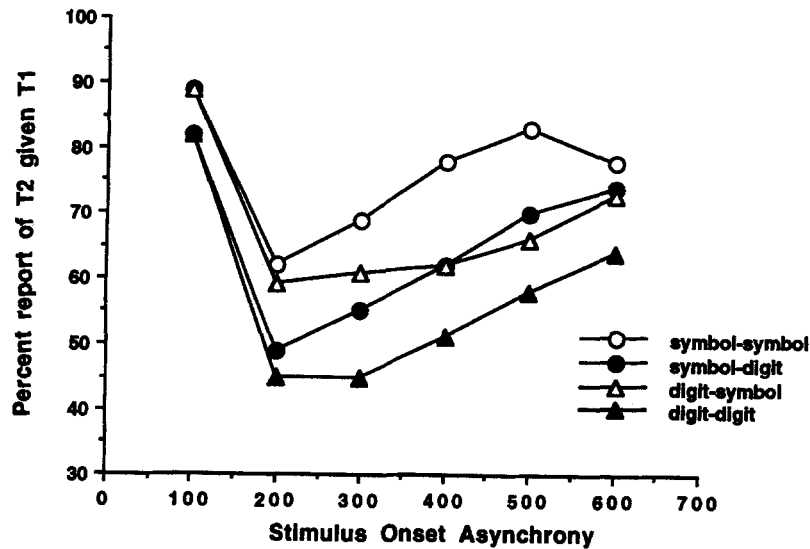
Figure 5. The effects of the type of distractor following the first target (T1) and the second target (T2) in Experiment 4.

when T2 immediately followed T1, T1 was correctly reported on 65% of the trials when D2 was a symbol and on 70% of the trials when D2 was a digit ($p > .13$).

This increase in performance for T1 when it was followed by a symbol rather than a digit supports our a priori expectation that letters would be easier to discriminate from these keyboard symbols than from digits. Of particular interest, however, is the effect that the type of D1 had on T2 identification. The main effect of D1 suggests that the AB deficit was modulated by the difficulty of T1 identification processing, as the two-stage model predicts.

However, while Experiment 4 highlights the local effects of the type of distractor immediately following T1 (as well as that following T2), it is clear from Figures 4 and 5 that an AB lag effect was still present, even when both T1 and T2 were followed by more readily discriminable symbol distractors. In Experiment 4, both letter targets still appeared among a set of mixed distractors, where roughly half of the distractors were less discriminable digits. This result raises the question of how T2|T1 performance is affected by the discriminability between the letter targets and the distractor set as a whole. In the following experiment, we varied the overall discriminability of the target and distractor sets. We predicted that an overall increase in discriminability would shorten the first stage of detection because viewers could lower their criterion for initiating Stage 2 target processing. As a consequence, Stage 2 processing would be completed more rapidly, thus reducing the size and duration of AB.

Experiment 5

In the previous experiment, we examined the effect of local discriminability using a mixture of digit and symbol distractors on each trial. Although main effects of local interference were found, an overall AB deficit persisted for

all conditions tested. Experiment 5 explored how overall discriminability between targets and distractors affects the magnitude and shape of the AB function. A separate question we also investigated was whether holding the serial position of T1 constant would reduce AB by allowing more rapid target detection. In this experiment, participants were divided into two groups. Each group participated in two blocked conditions. In the digit block, the task was identical to that of Experiment 1: Two letter targets appeared among digit distractors. In the symbol block, keyboard symbols were substituted for the eight digit distractors. For one group of participants, the serial position of T1 varied from trial to trial, as in Experiment 1 (each trial including 13 items). For the other group of participants, the procedure was simplified by presenting T1 in Serial Position 2 on all lists; the lag of T2 was varied as before (each trial included 9 items).

Method

The method was identical to that of the previous experiment, except as specified below.

Participants. Sixteen participants were tested in this experiment, 8 in each of the two groups (T1 fixed vs. T1 variable in serial position).

Design and procedure. There were two within-subject conditions. In the low-discriminability (digit) condition, the two target letters were embedded in an RSVP stream of digits. In the high-discriminability (symbol) condition, the letters were embedded in a stream of keyboard symbols. The two conditions were run in separate blocks: The order of the blocks was counterbalanced across participants in each group.

The letter targets and symbol and digit distractors were from the same sets used in Experiment 4. For one group of participants (the T1-variable group), each trial consisted of 13 items. The position of T1 was randomly permuted to appear an equal number of times
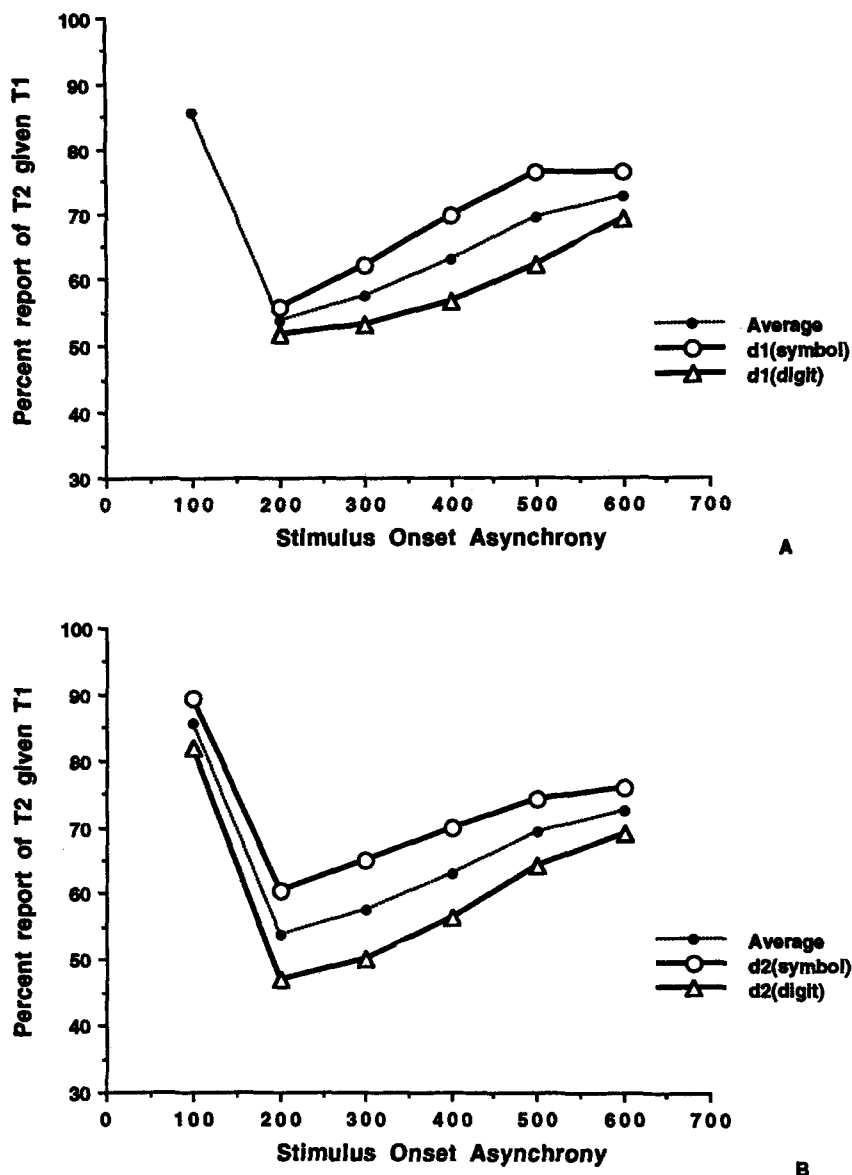
Figure 6. Effects in Experiment 4 of (a) the distractor following the first target (T1) and (b) the distractor following the second target (T2).

in Serial Positions 2–6. For the second group of participants (the T1-fixed group), each trial was shorter, consisting of 9 items. T1 always appeared in Serial Position 2 on the list. For both groups, T2 lag varied from Lag 1 to Lag 6. Thus, T2 never appeared as the last item in the list in either group. Each trial was preceded by a 300-ms fixation point (a plus symbol) and followed by an ampersand mask item. There were 20 trials at each of the six lag conditions, yielding a total of 120 trials for each block. Lag was randomized within each block.

The entire experiment was self-paced by the participant, who initiated each trial with a press of the spacebar on the keyboard. Participants were informed of the procedure and then were given a practice block of 20 trials using low-discriminability distractors (the digit set). No practice trials with the symbol distractors were given. Participants wrote their responses on a separate answer sheet.

Results and Discussion

Figure 7 shows the average percentage of trials in which T2 was correctly identified when T1 was correctly reported. The results clearly indicate a difference between the two distractor conditions, with digits producing lower T2 performance, $F(1, 14) = 80.4$, $p < .001$. There was a main effect of lag, $F(5, 70) = 18.2$, $p < .001$, and an interaction between distractor condition and lag, $F(5, 70) = 11.5$, $p < .001$. The main effect of T1 position variability, tested between groups, was not significant ($p > .55$). However, there was a three-way interaction between T1 position variability, distractor condition, and lag, $F(5, 70) = 2.7$, $p < .05$.
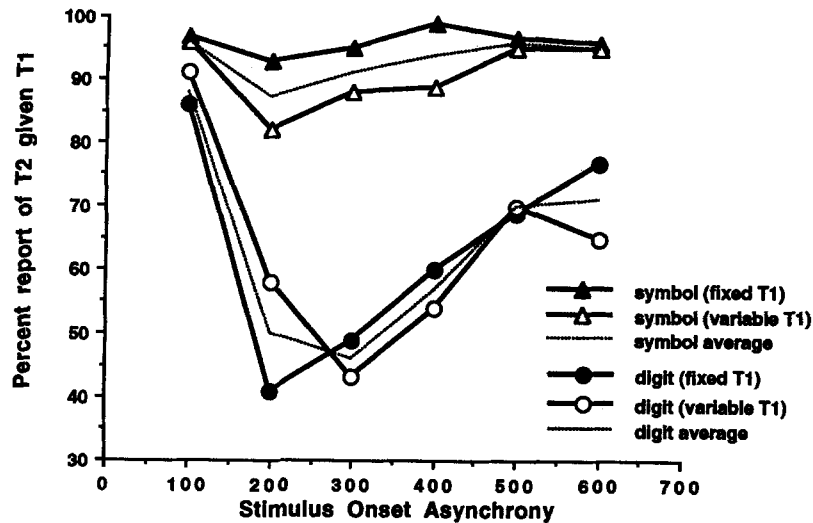
Figure 7. Effect of distractor set (symbol vs. digit) for fixed or varied serial position of the first target (T1) in Experiment 5.

Separate analyses were also carried out for each group of subjects (T1 variable and T1 fixed). In the T1-variable group, there was a main effect of distractor condition, $F(1, 7) = 79.5, p < .001$, a main effect of lag, $F(5, 35) = 10.9$, $p < .001$, and a significant interaction between condition and lag, $F(5, 35) = 6.0, p < .001$. A similar pattern of results was obtained for the T1 fixed group; there were main effects of distractor condition, $F(1, 7) = 30.3, p < .001$, and lag, $F(5, 35) = 8.7, p < .001$, and an interaction between distractor condition and lag, $F(5, 35) = 8.2, p < .001$.

Separate analyses of the two distractor conditions showed main effects of lag in the digit distractor condition, $F(5, 70) = 16.7, p < .001$, but no difference between the T1-fixed and the T1-variable groups, $F(1, 14) = 0.0, p > .99$, and no interaction, $F(5, 70) = 1.8, p > .11$. In the symbol distractor condition, there were main effects of lag, $F(5, 70) = 5.4, p < .001$, T1 variability, $F(1, 14) = 4.6, p < .05$, and their interaction, $F(5, 70) = 2.4, p < .05$. Inspection of the upper curves in Figure 7 indicated that there was little or no lag effect with symbol distractors when T1 was in a fixed serial position, but there was a small AB deficit when T1 was variable.

Thus, fixing the position of T1 and shortening the sequence to nine items resulted in improved performance, but only when the distractors were symbols. There was no consistent effect of fixing T1 when the distractors were digits. The AB effect was at least as large when T1 was fixed as when it was varied.

Excluding Lag 1, T1 was correctly identified on 98% of the trials in the symbol distractor condition and on 86% of the trials in the digit distractor condition, $F(1, 14) = 42.2, p < .001$. Thus, a letter target was easier to report when it appeared among discriminable symbols than among digits. It is surprising that there was no main effect of T1 predictability (T1 fixed = 93% and T1 variable = 91%) on T1 identification ($p > .41$) and no effect of T2 lag (excluding

Lag 1; $p > .65$). When T1 was immediately followed by T2, T1 was detected on 79% of the trials in the digit condition and on 86% of the trials in the symbol condition, $F(1, 14) = 4.2, p < .07$, and at this lag there was a main effect of T1 predictability, $F(1, 14) = 7.5, p < .05$, with T1 fixed reported on 88% of trials and T1 variable on 77%.

The main finding of the present experiment is that the deficit for detecting T2 was markedly attenuated when the targets were readily discriminable from the distractor set. Raymond et al. (1992; Experiment 3) have shown that when T1 is followed by a blank, the deficit is eliminated. This suggests that local interference by an immediately following item is a necessary condition for triggering the AB deficit. However, the near elimination of AB in the symbol condition here, together with the smaller reduction of AB in the D1 symbol condition of Experiment 4, suggests that the AB deficit is graded rather than triggered in an all-or-nothing fashion by any posttarget event. Although there was an effect of lag in the symbol condition, the AB effect was much weaker, and the significant interaction of lag and distractor condition indicates that the time course of AB was affected by the distractor condition. Moreover, the much greater reduction in AB when all the distractors were symbols, compared with the mixed-distractor conditions of Experiment 4, shows that it is not just D1 and D2 that determine the presence and time course of AB: The global distractor set also has a substantial effect. This result is consistent with our assumption that the global set influences the threshold criterion for detecting a target and initiating second stage processing.

Inversion Errors

In the experiments reported in this article, participants were encouraged but not required to report the target items

in the order that they appeared. The order of each participant's response (when two items were reported) was not recorded in Experiment 1. However, participants were observed to have made some order-inversion errors. In Experiment 2 and in all of the later experiments, the order in which participants reported the targets was recorded (for written responses, the left-hand item was assumed to be the first item reported). Trials in which both targets were reported were scored as correct, regardless of the order of report. In a separate analyses, all trials on which both T1 and T2 were reported were scored separately for each condition and SOA for the percentage of inversion errors. The means from Experiment 5 are shown in Figure 8.[2]

ANOVAs were carried out on the percentage of inversion errors at each SOA for the two distractor conditions (combining the two T1 position groups whose results were similar). There was a main effect of lag, $F(5, 70) = 13.2$, $p < .001$, with the majority of inversion errors occurring at Lag 1 (SOA = 100 ms). The results presented in Figure 8 indicate that the proportion of inversion errors decreased rapidly with increasing lag between T1 and T2. A significantly higher proportion of inversion errors was made in the more difficult high-similarity condition, $F(1, 14) = 12.7$, $p < .01$, suggesting that temporal uncertainty increases when the temporally adjacent letter targets are flanked by similar digit distractors. The interaction between condition and lag was also significant, $F(5, 70) = 5.7$, $p < .001$, and seems to be due to a floor effect in the symbol condition. The rarity of inversion errors at longer lags suggests that participants were following our instruction to try to report the two targets in the order in which they perceived them.

The occurrence of order errors at short target separations has been reported in many previous studies (e.g., Reeves & Sperling, 1986; Scarborough & Sternberg, 1967) and has been taken to indicate a failure of temporal resolution. The high proportion of 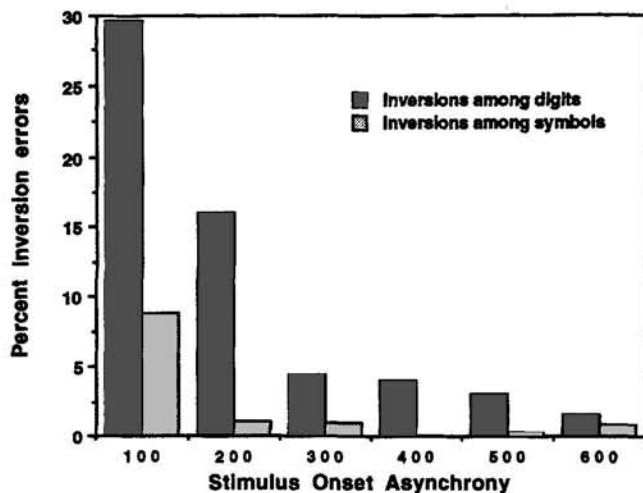inversion errors when the targets are adjacent is consistent with the hypothesis that second-stage processing of T1 overlaps with that of the following item. When the following item is a distractor, then the degree of visual similarity determines the amount of interference in the correct identification of T1. T1 was identified correctly on a higher percentage of trials when it was followed by a symbol than when it was followed by a digit. When the following item was another target, that saved T2 from the AB effect, at the cost of uncertainty in target order.

The prevalence of Lag 1 inversion errors in our task is analogous to the pattern of T1 misidentification errors reported in previous studies. When T1 is a colored letter among black letter distractors, the majority of identification errors involve intrusions rather than total misses (Raymond et al., 1992). A high proportion of these errors are +1 intrusions in which the letter distractor immediately following T1 is misreported as being the colored target (see also Broadbent & Broadbent, 1987; Gathercole & Broadbent, 1984; Lawrence, 1971; McLean et al., 1983).

In the two-stage model, we propose that the second, capacity-limited stage is initiated by detection of T1 but that it frequently includes both T1 and the following item. The evidence for frequency-inversion errors when T1 and T2 are adjacent is consistent with this hypothesis. That the inversion errors are markedly reduced (even at Lag 1) when the distractors are symbols supports the additional hypothesis that Stage 2 processing of T1 begins earlier and thus has less overlap with T2 at Lag 1 when targets are globally distinct from distractors.

Experiment 6

We examined the role of local discriminability in Experiments 3 and 4, and the effect of global discriminability in Experiment 5. The results indicated an effect of both local and global discriminability. In the next two experiments, we looked at the effects of local interference within a global set of digit distractors (Experiment 6) and symbol distractors (Experiment 7).

In Experiment 4, either a digit or a symbol immediately following T1 was sufficient to produce AB when the targets appeared among a mixed-distractor set of symbols and digits. In Experiment 6, we used a digit distractor set as our baseline and examined the effects of a single symbol distractor placed at either Lag 1 or Lag 2 in the RSVP stream. Experiment 6 was identical to Experiment 3, except that the blank intervals were replaced with an equals sign. In a pilot experiment in which 1 of 8 different symbols randomly replaced the blank, there was variation in the effect of different symbols. To reduce variability, we used only the equals sign in Experiment 6; it was the item that produced the smallest AB effect in the pilot study.

Figure 8. Percentage means of inversion errors in each condition, given report of both the first target (T1) and the second target (T2), in Experiment 5.

---

[2] Although not reported for the other experiments, the same pattern of inversion errors was observed in all of the AB experiments in this article (inversion data were not available for Experiment 1).

Method

The method was the same as that of Experiment 3, except as specified below.

Participants. There were eight participants in this experiment.

Design and procedure. The design and procedure were identical to those in Experiment 3, except that blanks were replaced with an equals sign. Thus, there were three conditions: no equals sign (only digit distractors), an equals sign at Lag 1, or an equals sign at Lag 2. All trials were intermixed randomly. The task was to report the two letters.

Results and Discussion

Figure 9 shows the mean percentage report of T2 given report of T1 for each condition. The results of the baseline condition are similar to the corresponding conditions in Experiments 1, 3, and 5, although the AB effect appears to be less marked in Experiment 6. Whether the equals sign appeared at Lag 1 or Lag 2, it had a similar interfering effect to that of a digit in the baseline condition. In a comparison of the Lag 1 symbol condition with baseline, there was a main effect of lag, $F(5, 35) = 10.7$, $p < .001$, but no effect of condition ($p > .26$) and no interaction between SOA and condition ($p > .14$). The apparent divergence of the two conditions was significant for Lag 5 (SOA = 500 ms), $t(7) = 2.2$, $p < .05$, one-tailed, suggesting quicker recovery in the Lag 1 symbol condition. The main result, however, was that the equals sign appearing immediately after T1 was sufficient to produce an AB deficit for T2. As expected, presenting the equals sign at Lag 2 also had no differential effect relative to the baseline. There was again a main effect of lag, $F(5, 35) = 11.7$, $p < .001$, but no interaction between condition and lag ($p > .23$).

The equals sign did affect T1 report, however. T1 was correctly reported on 87% of the trials in the baseline condition when it was immediately followed by a digit, while it was correctly reported on 96% of the trials when it was immediately followed by an equals sign, $F(1, 7) =$
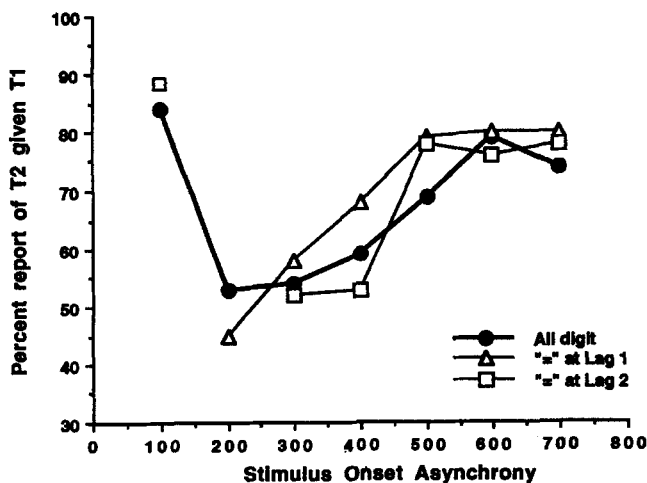

Figure 9. Effect of an equals sign among digits in Experiment 6.

37.3, $p < .001$. Thus, T1 was easier to identify and report when it was followed by a simple visual event than when it was followed by a digit. When T1 was followed by an equals sign at Lag 2, it was detected on 88% of the trials, compared with 87% in the baseline condition ($p = ns$). When T2 appeared at Lag 1, T1 was detected on 80% of the trials in the baseline condition and on 76% of the trials when T2 was followed by an equals sign.

The results are consistent with the results of Experiment 4 in showing that an easily discriminable, to-be-ignored visual event—that is, a symbol that was unlike both the letter targets and the digit distractors—is not equivalent to a blank interval. J. E. Raymond (personal communication, May 1992) and her colleagues obtained a similar finding using a simple dot pattern mask in their task in place of a distractor letter. Even though the equals sign interfered less with T1, it was sufficient to produce a salient AB effect, especially at the shorter lags.

Experiment 7

In the previous experiments, we found that the target report depended on the visual event directly following it. In Experiment 6, except for a single equals sign, T2 was embedded in a stream of digits, making it hard to identify. In Experiment 7, we switched the roles of symbols and digits, thus manipulating T1 identification difficulty while making T2 relatively easy to detect, since it appeared in a global environment of highly discriminable symbols, as in the symbol condition of Experiment 5.

Method

The method used was identical to that of Experiment 6, except as noted below.

Participants. Eight participants were tested in this experiment.

Design and procedure. The basic design of Experiment 7 was like that of Experiment 6, except where noted. Keyboard symbols replaced digit distractors and one (random) digit replaced the equals sign used in Experiment 6. Each trial in Experiment 7 consisted of nine items, with T1 always appearing as the second item in the list. The letter targets and digit and symbol distractors were from the same sets used in Experiment 4. The two target letters were randomly chosen from the target set for each trial, without replacement. Distractor items were also randomly chosen from their respective set, with the constraint that the same distractor item did not appear within the previous five serial positions in the RSVP list. In the all-symbol condition, all distractors were symbols. In the Lag 1 digit condition, a digit appeared immediately after T1. In the Lag 2 digit condition, a digit appeared at Lag 2, and either a target letter or a symbol appeared at Lag 1. Fifteen trials at each lag in each condition were run. All conditions were randomly intermixed, for a total of 240 trials, which were preceded by a block of 20 practice trials.

Results and Discussion

The results of the experiment are shown in Figure 10. In the all-symbol condition, there was an effect of lag, $F(5, 25) = 4.5$, $p < .01$. Inspection of the results in Figure 10
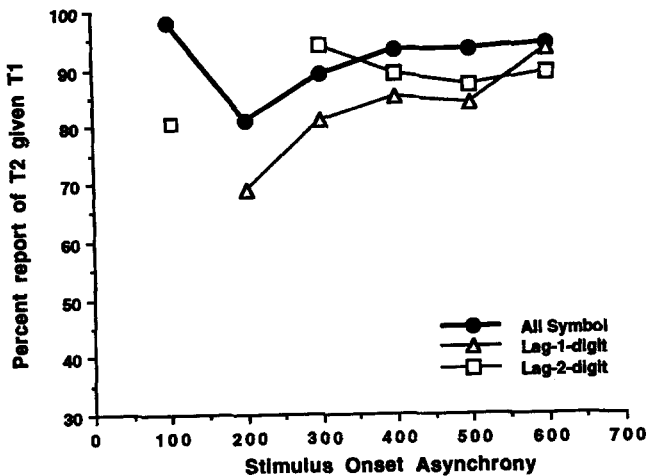
Figure 10. Effect of a digit among symbols in Experiment 7.

showed an AB effect that is much smaller than when the distractors were digits, although larger than the null effect in Experiment 5 with blocked all-symbol distractors. Comparing the baseline condition with the condition in which a digit distractor appeared immediately after T1, there was a main effect of condition, $F(1, 5) = 18.0, p < .01$, and lag, $F(4, 28) = 5.9, p < .01$, but the interaction was not significant ($p > .74$). Thus, the presence of a single digit immediately after T1 increased the AB deficit. When the single digit appeared at Lag 2, the pattern was quite different. There was no indication of an increased AB effect (compared with the all-symbol condition) after the digit (at SOAs of 300–600 ms; $p > .24$), although there was an interaction between condition and lag, $F(3, 21) = 3.8, p < .05$. More interesting was the highly significant difference between the baseline and Lag 2 digit condition at Lag 1, $t(7) = 5.5, p < .001$. It was more difficult to report the second of two adjacent letters when the second letter was immediately followed by a digit than when it was followed by a symbol distractor, which indicated that the item following T2 had an effect on its encoding, just as the item following T1 affected T1 encoding. The result is consistent with the results of Experiment 4. Excluding Lag 1, T1 was identified on 99% and 93% of the trials in the all-symbol and Lag 1 digit conditions, respectively, $F(1, 7) = 8.5, p < .05$. In the Lag 2 digit condition, T1 was identified on 98% of the trials.

In Experiment 7, both T1 and T2 were discriminable from all distractors, except for the occasional digit. Even when T1 identification was made difficult by an immediately following digit, AB was much reduced compared with AB in Experiment 4. The results from Experiments 6 and 7 suggest that the deficit was modulated by the amount of local interference with T1, but they also confirm a strong global effect of discriminability of the targets from the main distractor set.

General Discussion

The experiments reported in this article extend previous findings of a deficit for detecting the second of two targets appearing among distractors in RSVP (Broadbent & Broadbent, 1987; Raymond et al., 1992; Weichselgartner & Sperling, 1987). Previous studies had shown that the processing of T1 produces the deficit for T2. However, it was not clear whether the processing necessary for AB involved conjoining a target-defining feature with the correct features of the target (e.g., a white letter with that letter's identity) or whether registering any sort of target would be sufficient to produce AB under RSVP conditions. In the present experiments, targets were defined by their category: Participants were asked to detect and report letter targets among distractors presented for 100 ms per item. In Experiment 1, two letters were presented among digit distractors. When T1 was correctly reported, there was a strong deficit for reporting T2 appearing within an SOA of 200–400 ms, which demonstrated that an AB deficit occurred even for categorically defined targets appearing among distractors that were not members of the target set. This result suggests that the source of interference in AB is the attentional requirement of having to process a target that must be discriminated from a sequence of items.

In Experiment 2, we presented three targets and showed that the AB effect is produced anew by each target reported. We then examined the effect of discriminability in target search by manipulating both the overall discriminability between targets and distractors and the local discriminability between a target and the immediately following distractor. In Experiment 3, we replicated Raymond et al.'s (1992) results, showing that the AB deficit for T2 is highly reduced when the effective duration of T1 is extended by a blank frame. This result indicates that posttarget interference from an immediately following item is a necessary condition for AB. To evaluate the effect of distractor discriminability, we varied in Experiment 4 the type of distractor (digit vs. symbol) immediately following T1 and T2 in mixed sets of distractors. As predicted, T2|T1 performance was modulated not only by the type of distractor following T2, but also by the type of distractor following T1, suggesting that the AB deficit is a graded function of T1 processing difficulty.

To examine the effects of global target–distractor discriminability, we tested participants in separate blocks of all-digit distractors and all-symbol distractors in Experiment 5. Supporting the idea that target discriminability is a major factor, letter targets that appeared among discriminable symbols were much easier to detect than when they appeared among digit distractors. The effect of global discriminability was strong, in that the AB effect was virtually abolished when target letters appeared among symbol distractors. In the last two experiments, we manipulated local discriminability between T1 and the following distractor in a search task for letter targets among digit or symbol distractors (Experiments 6 and 7, respectively), with a single distractor from the other set on some trials. A difficult (but not an easy) local distractor following T1 again modulated the dominant effect of the global set of distractors.

In summary, the results indicate that target search performance in RSVP is a function of both global and local target–distractor discriminability and that the degree of

posttarget interference on T2|T1 is modulated by the difficulty of T1 processing. What is not clear from our experiments or from previous studies is exactly what features determine the discriminability between targets and distractors. Inspection of Figure 1 suggests that the symbols we used look less like letters than do the digits. As measured by the systematic relation between T1 performance and distractor type, letters and symbols were more readily discriminated in the present experiments than were letters and digits. Not only was there less visual overlap between the letters and symbols we used than between the letters and digits, it is also likely that the symbol set was less conceptually similar to letters than the digits were. Among other things, both digits and letters have names that are highly familiar, unlike most of the symbols. Thus, letters appearing among symbol distractors are presumably not only easier to process because of weaker visual-masking effects in our RSVP task, but are also probably more discriminable as conceptual entities. Further research is needed to clarify and separate the effects of visual similarity and categorical similarity.

A Two-Stage Model for the Attentional Blink

As discussed in the introduction of this article, Raymond et al. (1992) proposed that confusion between T1 and the immediately following item (specifically, confusion about how to conjoin the features of each item) triggers temporary attentional suppression for subsequent items. We propose instead that identification and consolidation of T1 is slowed down when there is an immediately following item, leading to a delay in allocating second-stage processing to T2 when it appears 200–400 ms after the onset of T1, which results in loss of T2 on some trials. While our experiments do not rule out an attentional suppression model, the pattern of results is more consistent with the proposed two-stage model. In the following sections, we discuss the explanatory power and generality of the model with respect to our own and other AB results, as well as to other paradigms in the selective attention literature.

The two-stage model we propose extends Broadbent and Broadbent's (1987) observations that early stages of detection are succeeded by more demanding and capacity-limited processes. This type of two-stage conceptualization dates back at least to Neisser's (1967) proposal that preattentive processes guide the operation of a focal attention stage. This proposal has been incorporated into various theories of spatial selective attention (e.g., Duncan, 1980, 1985; Hoffman, 1978; Treisman & Gelade, 1980; Wolfe, Cave, & Franzel, 1989) and of temporal selective attention (e.g., Gardner, 1973; Shiffrin & Gardner, 1972). The present two-stage model proposes that the AB deficit arises from a limited-capacity stage of processing and consolidation of the target after the target has been initially detected in the first stage.

First stage: Rapid detection. When items are presented for 100 ms each, virtually every item will be processed through the first stage in which features relevant for target detection are analyzed. Consistent with Broadbent and Broadbent (1987), we assume that processing at this stage of analysis allows for selection of candidate targets on the basis of independent feature cues, such as color or letter-case, or detection of targethood on the basis of categorical identity. On the basis of evidence from RSVP experiments involving word and picture stimuli, Potter (1976, 1983, 1984, 1993) proposed that short-lived conceptual representations are constructed for stimuli presented at rates as high as those used in AB studies (about 10 per second). The results of our experiments are consistent with this claim, suggesting that at presentation rates of 10 per second, the categorical identity of most of the items, and probably their specific identity (see Sperling et al., 1971), is briefly available and may serve as the basis of selection into subsequent stages. We assume that these initial representations are subject to rapid forgetting when there is interference from subsequent RSVP stimuli, unless they are selected for further processing and consolidation (Potter, 1976, 1993).

Second stage: Capacity-limited processing. In accord with Duncan (1980), we assume that the representations resulting from early levels of processing (first stage) cannot serve as the basis for subsequent report or response but require additional processing. A representation of a candidate target stimulus that is momentarily active must be transferred into a more durable representation (such as verbal short-term memory) to be available for subsequent report or, as Duncan suggests, even to serve as a basis for a manual response. This transfer requires second-stage processing which results in full identification and consolidation of the target for subsequent report. We consider this stage to be capacity-limited and to exceed the item's stimulus duration at the high presentation rates used in RSVP tasks. This stage of processing does not commence with the onset of a stimulus, but only after first-stage target detection. We hypothesize that the second stage is initiated by a transient attentional response that occurs on first-stage detection of a (probable) target. This attentional response actively selects and enhances processing of the target (Nakayama & Mackeben, 1989; Weichselgartner & Sperling, 1987). We also assume that the timing and resolution of this processing enhancement is such that T1 and the Lag 1 item (T2 or distractor) are likely to be processed together. Second-stage processing then identifies and consolidates the one or two targets present in the input, discarding distractor "noise." Until this second processing stage is completed, no subsequent items are processed beyond Stage 1. When T2 appears before the second stage is free, it will be detected by Stage 1 processing, but Stage 2 processing will be delayed. The longer the delay, the greater the probability that T2 will have been lost, according to our previous assumption that Stage 1 representations are short-lived.

T1 performance. Local discriminability between T1 and the following item affects the accuracy of T1 report. Effects of local interference on T1 detection were observed in all of the present experiments: a letter following T1 led to the least accurate T1 report, with T1 report improving when the following event was a digit, symbol, or blank (in that order). According to the model, the duration or efficiency of Stage

2 processing reflects T1 processing difficulty. In the following sections, we discuss how T1 processing difficulty correlates with report of T2.

T2 lag effect (Lag 1). If the item appearing at Lag 1 is processed together with T1, then when it is a target, it should benefit from Stage 2 processing. This prediction is supported by the lack of a deficit for T2 when it appeared at Lag 1 (Experiments 1–7). The high proportion of inversion errors at Lag 1 but not at other lags (Experiment 5) is consistent with the assumption that the two targets are processed together at Lag 1.

A further prediction is that T1 and T2 produce some mutual interference when both are processed together at Lag 1, and indeed T1 was least often reported when T2 immediately followed (see Table 1). Conversely, although T2 is more likely to be recalled at Lag 1 than at Lags 2–4 when T1 is reported, it is even more likely to be recalled if T1 is missed: For T2|T̄1 (T2 given that T1 was missed) in Experiment 1, the probability of recall was 94%, whereas for T2|T1 it was 76%. Similarly, the likelihood of reporting both T1 and T2 was lower at Lag 1 (51%) than at Lags 6–8 (71%) when the AB deficit was no longer in evidence (see Table 1). At Lag 1 only, T1 performance also appeared to be affected by the type of distractor that followed T2 (Experiments 3 and 6). When T2 was followed by a blank or symbol (presumably making T2 easier to process), there was more interference with T1 than when T2 was followed by another digit. Thus, T1 and T2 seem to be in competition when they are temporally adjacent, which is what the two-stage model predicts.

T2 lag effect (Lags 2–4). Resources from the capacity-limited second stage are unavailable for allocation to T2 while T1 is being processed. Thus, a second target appearing during this time has to wait. Within this interval, targets appearing at shorter SOAs must wait longer, on average, and are thus more susceptible to erasure by subsequent items. When the lag of T2 exceeds the duration of Stage 2 processing of T1, performance is no longer impaired.

Effects of T1 processing difficulty. If the duration of second-stage processing of a target is determined by the amount of interference from the following distractor (either visual feature overlap or conceptual similarity), then T2 performance should be affected. Experiments 3, 4, 6, and 7 all provide evidence of quicker recovery from the AB deficit as a function of the difficulty of discriminating between T1 and the following item. Differences in T1 processing difficulty did not always show up at Lag 2, but often only at later lags (3–5; see Experiments 4 and 6, Figures 5, 6A, and 9). This result is expected if we assume that the second stage of processing overlaps with Lag 2, even when the distractors are easy to discriminate. The difference between easy and difficult distractors only begins to show up at Lag 3.

Global distractors and target criterion. Even when the items following T1 and T2 were held constant, the global distractor set had an influence on performance (Experiments 4–7). We hypothesize that participants establish a target detection criterion in accordance with the set of distractors, lowering the criterion when discrimination between the target and distractors is easy (as with symbol distractors) and raising it when distractors are more difficult to discriminate from targets (as with digit distractors). This effect is predicted to be largely independent of the effect of the specific distractor following a target.

A lower criterion means that fewer target-relevant features must be detected before Stage 2 processing is triggered. The earlier that second-stage processing begins, the sooner it will be completed, the smaller the likelihood that the Lag 1 item will be included with the target, and the less overlap in processing there will be with subsequent items. When targets and distractors are highly discriminable in a given block (Experiment 5, symbol condition), the criterion can be set very low. The minimal AB deficit and the lower proportion of Lag 1 inversion errors in the symbol condition support this suggestion. Evidently, Stage 2 processing was initiated before the following item was available to be included in processing and was usually completed before T2 was lost through interference. A comparison across experiments showed that a small deficit occurred on all-symbol trials when they were intermixed with trials that included a single digit distractor (Experiment 7) or when all-symbol trials were blocked but the position of T1 was randomized (Experiment 5). The criterion for target detection can presumably be set lower when the all-symbol trials are blocked and when T1 position is fixed.

Comparing the Two-Stage Model With Previous AB Models

The present model ties together several ideas brought up by previous models of AB. As discussed before, Broadbent and Broadbent (1987) were the first to suggest that the AB deficit might be indicative of a two-stage process of initial detection of target features followed by more demanding processes of identification. However, their initial proposal did not extend much beyond this general observation and is not specific enough to account either for the lack of AB at Lag 1 or for the effect of local and global distractor interference. Weichselgartner & Sperling (1987) proposed a two-component attentional mechanism to account for post-T1 deficits, and according to this model, T2 appearing at Lag 1 is picked up along with T1 by virtue of a transient attentional response. However, in their model report of later targets is processed by a second, qualitatively different, sustained component of attention. In a task involving three targets (Experiment 2), we showed that the AB pattern (Lag 1 benefit, Lags 2–4 deficit) iterates for each occurrence of a correctly reported target, a result that would not be expected if T2 were detected by a sustained component. Weichselgartner & Sperling's model also makes no prediction about variation in the type of distractor following T1.

Raymond et al. (1992) reported several results important to understanding the cause of AB. Of particular importance is the finding that the item appearing immediately after T1 plays a critical role in producing AB. Thus, it is not target identification alone, but interference from immediately fol-

lowing visual stimuli, that produces the deficit for T2. The role of this critical item immediately following T1 (D1) has been a major focus of this article.

Our experiments and two-stage model extend the results reported by Raymond et al. (1992) and share several assumptions about the AB effect, especially in regard to the role of D1 in the AB effect. Both models also assume that T1 and D1 are usually processed together. Furthermore, Raymond et al. (1992) suggested that the amount of interference is proportional to the degree of similarity between D1 and T1 (Duncan & Humphreys, 1989). The lack of AB at Lag 1 and the high proportion of inversion errors for T1 and T2 at Lag 1 shown in our results are consistent with the assumption that T1 and D1 are processed together. The effect of local discriminability between targets and their immediately following distractors, as well as the effect of the global distractor set, supports the predicted effect of greater interference as a function of increasing similarity.

However, there are major differences between our two-stage model and the attentional suppression model proposed by Raymond et al. (1992). First, they proposed that the main source of interference with T1 identification is a feature-conjunction confusion arising when the target-defining feature (a white letter) has to be conjoined with letter features and when both T1 and the Lag 1 letter are present in a sensory store that does not represent serial order. They proposed that when the system detects the potential for confusion, a suppressive mechanism is initiated to eliminate further confusion. Our results showing AB with categorically defined targets suggest that the deficit is not the result of having to conjoin a target-signaling feature with target identity.

As discussed earlier, Raymond et al. (1992) explained the AB deficit as the result of an inhibitory attentional suppression mechanism, which is invoked by the visual system to prevent further interference from items following T1. In particular, this inhibitory process is described in terms of an attentional gate that becomes both shut and locked in the presence of confusion, thus producing an attentional blink. Because their T2 task (detecting the presence or absence of the letter X) was assumed to require only a minimal level of processing, Raymond et al. concluded that the suppression mechanism blocks off further visual processing at a relatively early stage. However, detection of a letter X appearing among a rapid sequence of other distractor letters may be a more demanding process than it seems. Yes–no detection of even a prespecified letter appearing among other letters may require identification of the features of each distractor letter (see Sperling et al., 1971). The main point is that if their T2 probe task was actually more complex than Raymond et al. assumed, there is no reason to conclude that visual processing is suppressed at an early level in AB. Instead, Stage 1 processing of subsequent items may continue, as we propose.

While the suppression model takes into account the effect of local similarity in modulating the size of the AB deficit, it is not clear how this fits with a shut-and-locked attentional mechanism. To explain the effects of local interference on the size of the AB deficit, as found in the present study, the suppression model must postulate that increased local interference results in either (a) increased probability of a shut-and-locked mechanism being activated or (b) stronger inhibition of subsequent stimuli. The latter account requires an added mechanism; as it was described, the locking mechanism is ballistic in that, once initiated, it is insensitive to posttarget events. The main effects of D2 immediately following T2 in Experiment 4 suggest that the detectability of T2, not just that of T1, must be taken into consideration.

In short, although the AB effect can be explained by the operation of the hypothesized attentional gate, some further specification of the operational characteristics of the attentional gate is needed. Our model provides such an account. Rather than hypothesizing that AB is due to some inhibitory mechanism, we propose that the AB deficit reflects a fundamental characteristic of visual processing: that there is a limited-capacity stage for identifying and consolidating targets in RSVP, with characteristics that we have outlined here.

The Two-Stage Model and Other Target-Interference Paradigms

It has been well established in a variety of other target search paradigms that simultaneous targets in a single array produce mutual interference. Duncan (1980) reviewed the evidence and presented several experiments to support his proposal for two levels of perceptual representation. At the first, preattentive level, targets and nontargets may be distinguished in parallel (Duncan used categorically defined alphanumeric targets and distractors). However, targets must pass through a limited-capacity system to a second level before forming a reportable perception. Multiple targets compete for access through the limited-capacity system, resulting in impaired detection performance for any given target when multiple targets are presented simultaneously. Our two-stage model for target detection in RSVP is consistent with Duncan's work, if one assumes that the onset of limited-capacity processing is affected by the observer's global detection criterion and that its duration and accuracy are affected by the following distractor. When targets are presented in RSVP, the AB deficit shows how this mutual target interference extends over time.

Duncan and Humphreys (1989) proposed a comprehensive theory based on the effects of similarity to explain a wide range of experimental results from the spatial visual-search literature. The three main components of Duncan and Humphreys's model are (a) a parallel stage of perceptual description that produces a structured representation of items in the visual field at a number of spatial scales, (b) the matching of these representations against an internal template of the target, and (c) the selection of template-matching items for entry into visual short-term memory, which is limited in capacity. According to the model, access to visual short-term memory is strictly limited, and items are competitively assigned weightings according to their degree of match to current target templates. Thus, target search difficulty in spatial displays increases as a function of

increased similarity of targets to nontargets and decreased similarity between nontargets. Though the similarity-based model for visual search is specific to simultaneous arrays, not RSVP target search, our experimental results and the two-stage search model we propose are consistent with certain aspects of Duncan and Humphreys's model, especially with regard to the predicted effects of target–distractor similarity. On the other hand, Duncan and Humphreys's model does not address lag effects or multiple targets and the AB effect.[3]

A related paradigm for studying processing limitations is the overlapping tasks paradigm (otherwise known as the psychological refractory period, PRP, paradigm) (Pashler, 1984, 1992; Welford, 1952). Here, two stimuli (S1 and S2) are presented in a sequence, and participants are asked to make a separate response (R1 and R2) to each stimulus in the order of presentation. As the SOA is decreased between S1 and S2, such that S2 is presented before R1 is made, interference occurs between the two tasks, which results in increased reaction time for R2. One class of models to explain the slowing of R2 proposes that a bottleneck occurs, such that certain stages of processing cannot be performed simultaneously on more than one input (McCann & Johnston, 1992; Pashler, 1984; Pashler & Johnston, 1989; Welford, 1952). Thus, S2 processing that involves this stage must be postponed until response selection for S1 is completed. Another class of models to explain the PRP effect does not posit a bottleneck but instead proposes that many cognitive operations drawn on a common pool of resources (Kahneman, 1973). When multiple stimuli compete for the limited resources in this pool, the processing of some or all of the stimuli is slowed or degraded. It is difficult to distinguish empirically between these two models. However, there is one pattern in our results that is more readily explained by a serial bottleneck than by a parallel competition for resources: the delay of onset of differential effects of the distractor following T1, on T2 report (see Figures 5 and 6). (This pattern has also been seen in experiments not reported here.) Such a pattern is consistent with a bottleneck that persists until at least Lag 3, even in the easier of two conditions. The pattern is not easily explained by a competing-resources model.

An important issue in understanding the PRP effect is the identity of the cognitive process that causes the delay of R2. The stages of processing of S1 and S2 can be roughly divided into three components: perception, response selection, and response execution. The evidence suggests that it is response selection that constitutes a bottleneck in PRP dual tasks (see Pashler, 1984, 1989; Pashler & Johnston, 1989).

It is not yet clear whether the response selection bottleneck for performing speeded vocal or manual responses in the PRP paradigm is the same bottleneck we propose as the source of AB. For an auditory first task (S1) and a visual second task (S2, report the highest digit among an array of digits), Pashler (1989) showed that perceptual processing of a masked visual display (S2) can proceed while response selection for S2 is underway. Only the corresponding response selection stage for S2 is delayed, which results in R2 slowing (but no increase in errors) when a speeded response

is required. However, when S1 is also a visual task (orientation feature search), then there is a dramatic increase in R2 errors when the SOA is short. Together, these results suggest that the source of interference producing RT slowing in PRP tasks may be different from the error-producing interference that occurs between two visual tasks. The limited-capacity process producing AB probably falls under the latter category; in that case, the bottleneck in the present model differs from the bottleneck causing PRP slowing. However, more directly comparable experiments will be required to clarify the relation between the two paradigms and their models.

Another deficit that appears in RSVP tasks and shares similarities with the AB deficit is repetition blindness (RB; Bavelier, 1992; Bavelier & Potter, 1992; Kanwisher, 1987, 1991; Kanwisher & Potter, 1989, 1990). Whereas AB is a deficit in reporting a second RSVP target when T1 and T2 are different, RB is the additional difficulty of reporting T2 when it happens to be the same stimulus as T1 (e.g., both are the letter *A*). Unlike the AB paradigm, however, in most RB studies participants are required to report all the items presented, not just the two critical items. The presentation rates used are comparable with those in AB studies, and the lags at which RB is found are similar. RB has typically been studied with RSVP sentences in which AB does not generally occur, although both AB and RB occur for unrelated stimuli such as random letters (for RB, see Kanwisher, 1991, and Bavelier & Potter, 1992; for RB and AB, see Chun & Potter, 1992). Also, under certain conditions AB occurs without RB (Chun & Potter, 1993; Ward, Duncan, & Shapiro, 1992), suggesting that the two deficits, though closely related, are doubly dissociable and reflect independent limitations in visual processing. Nonetheless, the effect of a repetition is additive with AB: If T2 survives AB, it is at risk for RB when T1 and T2 are identical. Apparently, RB results because T2 is perceived as the same event as T1 and so fails to be registered independently.

A question not addressed by the present experiments is whether the selection of a subset of targets from an RSVP stream involves additional processes that would not be present in whole report of every item in the sequence. AB studies use a partial report task, in which targets need to be selected from the RSVP stream for subsequent report; an additional cost may exist for such monitoring and selection. Unpublished pilot experiments from this lab that compare whole and selective report using letter and digit stimuli in short RSVP sequences suggest that a whole report strategy does not increase the absolute probability of report of an

---

[3] Shapiro and Raymond (1994) proposed a reformulation of their attentional suppression model to a similarity model based on Duncan and Humphreys (1989). This model is also substantially different from the present two-stage model. Duncan, Ward, and Shapiro (1994) suggested that "attentional dwell time" has a standard duration of 300–500 ms, with many items processed in parallel, but they failed to give a satisfactory account of the processing of nontargets and targets in an RSVP stream, in which the attentional stage is not triggered until a target is preattentively detected. A two-stage model like that we propose gives a better account of such sequential processing.

item in T2 position. However, the number of items to be reported was confounded with whole report versus selection, making it difficult to reach a firm conclusion about a possible selection cost in partial report.

The issue of selection versus whole report may be relevant to understanding the lack of intermittency on individual trials in perceiving and recalling RSVP sentences (e.g., Potter, 1984; Potter, Kroll, Yachzel, Carpenter, & Sherman, 1986). Studies involving words that constitute contextually meaningful and grammatical sentences are typically whole report tasks, and these tasks can be performed with high accuracy at the same presentation rates used in AB studies. Thus, it is not clear from existing research whether the apparent absence of AB-like intermittency in whole report of sentences indicates that there is an additional selection cost in the AB paradigm, or whether the speed of encoding successive words in meaningful sentences avoids the AB deficit.

The lack of AB for RSVP sentences presented at the same rate that produces AB in target search does, however, lend support to the idea that the AB deficit is not simply an effect of low-level masking, but rather reflects a bottleneck in higher level processing. According to our two-stage model, AB occurs in RSVP target-search tasks because of a rate-limiting stage between initial perception and subsequent report. Thus, our model can easily account for the lack of AB for sentences by assuming that contextual cues and syntactic parsing mechanisms allow for rapid and efficient encoding and retention of the incoming word stream. Support for this assumption can be found in a study by Forster (1970), who presented lists of six words at 62.5 ms per word. Performance for a string of words was highest when the words formed a simple sentence, lower for complex sentences, and significantly worse for scrambled strings. Pfafflin (1974) and Potter et al. (1986) also have reported that recall performance for scrambled sentences is much poorer than for the same sentence presented in normal order.

In summary, the present results support a two-stage model in which a first stage of perceptual detection is followed by a capacity-limited second stage that is required for the conscious retention of a target's identity for subsequent report. This stage constitutes a bottleneck that produces a transient deficit for second-stage processing of subsequent targets. This deficit extends over a period of up to 500 ms after onset of T1, during which performance on T2 improves with increasing lag (between Lags 2 and 5). The amount of interference from T1 is modulated by the difficulty of discriminability between T1 and the following item, as well as by the global discriminability between targets and distractors, and is largely eliminated by a blank interval immediately after T1 (see also Raymond et al., 1992). The lack of a deficit for a second target appearing at Lag 1 and the high proportion of order inversion errors made in that condition are consistent with the model's assumption that T1 and an immediately following item (whether another target or a distractor) are more likely to be processed together in Stage 2.

Although the present study focused on temporal limitations of search for multiple targets in an RSVP display, the model we propose is consistent with an array of models proposed for spatial visual search, and further research may allow generalization to the wide range of tasks in which attention must be deployed among multiple objects across space and over time.

References

Bavelier, D. (1992). Phonological repetition blindness. Unpublished doctoral dissertation, Massachusetts Institute of Technology, Cambridge.

Bavelier, D., & Potter, M. C. (1992). Visual and phonological codes in repetition blindness. Journal of Experimental Psychology: Human Perception and Performance, 18, 134–147.

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

Chun, M. M., & Potter, M. C. (1992). Interference in detecting multiple RSVP targets: Effects of similarity. Poster session presented at the 1992 annual meeting of the Psychonomics Society, St. Louis, MO.

Chun, M. M., & Potter, M. C. (1993). Interference in detecting multiple targets in a sequence: A dissociation between the attentional blink and repetition blindness. (Abstract No. 2592). Sarasota, FL: Investigative Ophthalmology and Visual Science.

Duncan, J. (1980). The locus of interference in the perception of simultaneous stimuli. Psychological Review, 87, 272–300.

Duncan, J. (1983). Perceptual selection based on alphanumeric class: Evidence from partial reports. Perception & Psychophysics, 33, 533–547.

Duncan, J. (1985). Visual search and visual attention. In M. I. Posner & O. S. M. Marin (Eds.), Attention and Performance (Vol. 11, pp. 85–105). Hillsdale, NJ: Erlbaum.

Duncan, J., & Humphreys, G. (1989). Visual search and stimulus similarity. Psychological Review, 96, 433–458.

Duncan, J., Ward, R., & Shapiro, K. (1994, May 26). Direct measurement of attentional dwell time in human vision. Nature, 369, 313–315.

Forster, K. I. (1970). Visual perception of rapidly presented word sequences of varying complexity. Perception & Psychophysics, 8, 215–221.

Gardner, G. T. (1973). Evidence for independent parallel channels in tachistoscopic presentation. Cognitive Psychology, 4, 130–155.

Gathercole, S. E., & Broadbent, D. E. (1984). Combining attributes in specified and categorized target search: Further evidence for strategic differences. Memory & Cognition, 12, 329–337.

Hoffman, J. E. (1978). Search through a sequentially presented visual display. Perception & Psychophysics, 23, 1–11.

Horlitz, K. L., Johnston, J. C., & Remington, R. W. (1992). What makes it difficult to process multiple targets in rapid serial visual presentation (RSVP)? Paper presented at the 1992 annual meeting of the Psychonomics Society, St. Louis, MO.

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

Kahneman, D. (1973). Attention and effort. New York: Prentice Hall.

Kanwisher, N. G. (1987). Repetition blindness: Type recognition without token individuation. Cognition, 27, 117–143.

Kanwisher, N. (1991). Repetition blindness and illusory conjunctions: Errors in binding visual types with visual tokens. Journal of Experimental Psychology: Human Perception and Performance, 17, 404–421.

Kanwisher, N., & Potter, M. C. (1989). Repetition blindness: The effects of stimulus modality and spatial displacement. *Memory & Cognition, 17,* 117–124.

Kanwisher, N. G., & Potter, M. C. (1990). Repetition blindness: Levels of processing. *Journal of Experimental Psychology: Human Perception and Performance, 16,* 30–47.

Krueger, L. E. (1984). The category effect in visual search depends on physical rather than conceptual differences. *Perception & Psychophysics, 35,* 558–564.

Lawrence, D. H. (1971). Two studies of visual search for word targets with controlled rates of presentation. *Perception & Psychophysics, 10,* 85–89.

McCann, R. S., & Johnston, J. C. (1992). Locus of the single-channel bottleneck in dual-task interference. *Journal of Experimental Psychology: Human Perception and Performance, 18,* 471–484.

McLean, J. P., Broadbent, D. E., & Broadbent, M. H. P. (1983). Combining attributes in rapid serial visual presentation tasks. *Quarterly Journal of Experimental Psychology, 35A,* 171–186.

Nakayama, K., & Mackeben, M. (1989). Sustained and transient components of focal visual attention. *Vision Research, 29,* 1631–1647.

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

Pashler, H. (1984). Processing stages in overlapping tasks: Evidence for a central bottleneck. *Journal of Experimental Psychology: Human Perception and Performance, 10,* 358–377.

Pashler, H. (1989). Dissociations and dependencies between speed and accuracy: Evidence for a two-component theory of divided attention in simple tasks. *Cognitive Psychology, 21,* 469–514.

Pashler, H. (1992). Attentional limitations in doing two tasks at the same time. *Current Directions in Psychological Science, 1,* 44–48.

Pashler, H., & Johnston, J. C. (1989). Chronometric evidence for central postponement in temporally overlapping tasks. *The Quarterly Journal of Experimental Psychology, 41A,* 19–45.

Pfafflin, S. M. (1974). The total time hypothesis, recall strategies, and memory for rapidly presented word strings. *Memory & Cognition, 2,* 236–240.

Potter, M. C. (1976). Short-term conceptual memory for pictures. *Journal of Experimental Psychology: Human Learning and Memory, 2,* 509–522.

Potter, M. C. (1983). Representational buffers: The eye–mind hypothesis in picture perception, reading and visual search. In K. Rayner (Ed.), *Eye movements in reading: Perceptual and language processes* (pp. 423–437). New York: Academic Press.

Potter, M. C. (1984). Rapid serial visual presentation (RSVP): A method for studying language processing. In D. E. Kieras & M. A. Just (Eds.), *New methods in reading comprehension research* (pp. 91–118). Hillsdale, NJ: Erlbaum.

Potter, M. C. (1993). Very short-term conceptual memory. *Memory & Cognition, 21,* 156–161.

Potter, M. C., Kroll, J. F., Yachzel, B., Carpenter, E., & Sherman, J. (1986). Pictures in sentences: Understanding without words. *Journal of Experimental Psychology: General, 115,* 281–294.

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

Reeves, A., & Sperling, G. (1986). Attention gating in short-term visual memory. *Psychological Review, 93,* 180–206.

Scarborough, D. L., & Sternberg, S. (1967). *Processing items and their order in sequential visual displays.* Paper presented at the 1967 annual meeting of the Eastern Psychological Association, Boston, MA.

Shapiro, K. L., & Raymond, J. E. (1994). Temporal allocation of visual attention: Inhibition or interference? In D. Dagenbach & T. H. Carr (Eds.), *Inhibitory mechanisms in attention, memory and language* (pp. 151–188). Boston: Academic Press.

Shiffrin, R. M., & Gardner, G. T. (1972). Visual processing capacity and attentional control. *Journal of Experimental Psychology, 93,* 72–82.

Sperling, G., Budiansky, J., Spivak, J. G., & Johnston, M. C. (1971). Extremely rapid visual search: The maximum rate of scanning letters for the presence of a numeral. *Science, 174,* 307–311.

Treisman, A., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology, 12,* 97–136.

Volkman, F. C., Riggs, L. A., & Moore, R. K. (1980). Eyeblinks and visual suppression. *Science, 207,* 900–922.

Ward, R., Duncan, J., & Shapiro, K. (1992). *The attentional blink does not require selection among a nontarget stream.* Paper presented at the 1992 annual meeting of the Psychonomics Society, St. Louis, MO.

Weichselgartner, E., & Sperling, G. (1987). Dynamics of automatic and controlled visual attention. *Science, 238,* 778–780.

Welford, A. T. (1952). The "psychological refractory period" and the timing of high-speed performance: A review and theory. *British Journal of Psychology, 43,* 2–19.

Wolfe, J. M., Cave, K. R., & Franzel, S. L. (1989). Guided search: An alternative to the feature integration model for visual search. *Journal of Experimental Psychology: Human Perception and Performance, 15,* 419–433.