Psychological Research (2005) 69: 191–200
DOI 10.1007/s00426-004-0173-x

ORIGINAL ARTICLE

Vincent Di Lollo · Jun-ichiro Kawahara
S.M. Shahab Ghorashi · James T. Enns

The attentional blink: Resource depletion or temporary loss of control?

Received: 15 May 2003 / Accepted: 13 January 2004 / Published online: 29 April 2004

Abstract Identification of the second of two targets is impaired if it is presented less than about 500 ms after the first. Theoretical accounts of this second-target deficit, known as *attentional blink* (AB), have relied on some form of limited attentional resource that is allocated to the leading target at the expense of the trailing target. Three experiments in the present study reveal a failure of resource-limitation accounts to explain why the AB is absent when the targets consist of a stream of three items belonging to the same category (e.g., letters or digits). The AB is reinstated, however, if an item from a different category is inserted in the target string. This result, and all major results in the AB literature, is explained by the hypothesis that the AB arises from a temporary loss of control over the prevailing attentional set. This lapse in control renders the observer vulnerable to an exogenously-triggered switch in attentional set.

Introduction

Attentional limitations in visual information processing are in evidence when two target stimuli must be identified in rapid sequence. Accuracy is nearly perfect for the first target but is dramatically reduced for the second (Raymond, Shapiro, & Arnell, 1992; Ward, Duncan, & Shapiro, 1996). Much of the evidence regarding this second-target deficit, also known as the *attentional blink* (AB), comes from studies in which two targets are inserted in a stream of distractors displayed in rapid serial visual presentation (RSVP). To study the temporal course of the second-target deficit, the lag between the two targets is varied systematically in steps of about 100 ms. The deficit is commonly found to be largest when the inter-target lag is short, and to diminish progressively as the lag is increased.

V. Di Lollo (✉) · S.M. Shahab Ghorashi · J. T. Enns
Department of Psychology, University of British Columbia,
2136 West Mall, Vancouver, B.C. V6T 1Z4, Canada
E-mail: vince.dilollo@ubc.ca

J. Kawahara
Hiroshima University, Hiroshima, Japan

Theoretical accounts of the AB have relied on some form of limited attentional resource that is allocated to the leading target at the expense of the trailing target. For example, the *attentional dwell-time hypothesis* holds that the AB occurs because resources required commonly by the two targets are unavailable for the second target while the first target is being processed (Ward et al., 1996). Similar assumptions underlie the *two-stage model* of Chun and Potter (1995) in which the AB deficit is said to occur when the second target arrives while the resources at a high-level stage are preempted by the first target. Jolicoeur and Dell'Acqua's (1998) *PRP model* is similarly predicated on resource limitations leading to a bottleneck at a late processing stage. Finally, resource depletion underlies the *interference model* of Shapiro, Raymond, and Arnell (1994) in which resources are said to be allocated in large part to the first target, and in diminishing amounts to the ensuing items in the RSVP stream.

A common theme in these accounts is the depletion of a limited resource by the first target with a consequent second-target deficit. Accordingly, the deficit is expected to be greatest when the second target is presented directly after the first, namely, at the shortest inter-target lag known as *Lag 1*. The deficit is then expected to diminish progressively at longer lags, reflecting the increased availability of resources for the second target as processing of the first target nears completion.

Precisely such monotonic trends over lags have been found in about half the experiments reported in the literature (Visser, Bischof, & Di Lollo, 1999). The remaining experiments, however, yielded non-monotonic *U*-shaped trends, revealing the phenomenon of *Lag 1 sparing*: Identification of the second target was virtually unimpaired at Lag 1, but dropped dramatically at Lag 2 before recovering at longer lags. It goes without saying that limited-resource models can account for the monotonic trend but not for the non-monotonic trend illustrated by Lag 1 sparing. We believe this to be a critical failure of limited-resource models.

To account for Lag 1 sparing, resource-depletion models have postulated an attentional gate that opens rapidly upon arrival of the first target but closes sluggishly, thus allowing the next item in the stream (i.e., the Lag 1 item) to gain access to processing resources along with the first target. If the Lag 1 item happens to be the second target, both targets are processed concurrently, and the AB deficit is avoided.

But there are problems with this account. The concept of a sluggish attentional gate is clearly ad hoc, in that it lies outside the conceptual framework of the resource-depletion hypothesis. Furthermore, the hypothesis is internally inconsistent as to the limit beyond which resources begin to deplete. In studies that exhibit Lag 1 sparing, the limit is set at two items (the first target and the Lag 1 item); in studies that do not exhibit Lag 1 sparing, however, the limit is only one item (the first target). But these estimates fall far short of the well-established limit of over four items, even when presented simultaneously in a brief display followed by a mask (Sperling, 1960). Finally, the account fails to explain why the second target can pass through the gate in about half of the experiments but not in the other half. The gating account is, therefore, insufficient. Factors other than a sluggish attentional gate must also be at work.

Here, we propose an alternative account based not on depletion of a limited resource but on a temporary loss of perceptual control that occurs during the process of target identification. Attentional switching has been implicated as a determining factor in recent AB studies (Kawahara, Zuvic, Enns, & Di Lollo, 2003; Potter, Chun, Banks, & Muckenhoupt, 1998; Visser et al., 1999). It has been suggested that attentional switching involves a reconfiguration of the system aimed at handling the incoming stimuli with maximum efficiency (Meiran, 1996; Monsell, 1996; Visser et al., 1999). Reconfiguration is part of a comprehensive goal-directed process aimed at tuning the visual system to those attributes and characteristics of incoming stimuli that are likely to prove useful for performing the task at hand. Monsell (1996) has referred to this process as *task-set reconfiguration*; William James (1890/1950) termed it *ideational preparation* or *adaptation of attention*.

Given that two targets are presented sequentially in AB experiments, it is reasonable to assume that the system is initially configured to optimize performance on the first target, which is therefore processed quickly and accurately. We can also assume that the initial configuration is governed endogenously by means of signals from higher brain regions, such as the prefrontal cortex, that function as a general-purpose central processor (e.g., Goldman-Rakic, 1987). In practice, this means that, at the start of an RSVP sequence, the system is configured as an input filter tuned to exclude the leading distractor items and to pass the first target. Filters such as these were popularized in the selective-attention literature of the 1950s and 1960s (Broadbent, 1958; Cherry, 1953; Treisman, 1960). The present filter is assumed to involve multiple regions of the visual system, and to be volatile rather than static. Namely, in order to remain viable, the filter is held to be in need of periodic maintenance signals from the central processor. A similar notion of volatile attentional filtering was proposed by William James (1890/1950) as follows: "*There is no such thing as voluntary attention sustained for more than a few seconds at a time*. What is called sustained voluntary attention is a repetition of successive efforts which bring back the topic to the mind" (p. 420; italics in the original).

We make the critical assumption that, while the central processor continues to issue appropriate control signals, the system's configuration cannot be altered exogenously by the leading distractors in the RSVP stream. Upon arrival of the first target, however, the central processor becomes engaged in stimulus processing and response planning, and can no longer issue the required control signals. This concept of an interruptible control mechanism dedicated to a single task is in keeping with such concepts as Shallice's (1994) *supervisory attentional system*, Shiffrin and Schneider's (1977) *controlled processing*, or Baddeley's (1986) *central executive*. In simple terms, an interruption in the functioning of the control mechanism corresponds to a lapse of attention.

What happens during such a lapse is critical for our account of the second-target deficit. We propose that, in the absence of endogenous control signals, the system's configuration comes under exogenous control by the ensuing item, namely, the item in the Lag 1 position. If that item belongs to the same category as the first target, it will match the current configuration of the system and will gain access to further processing, leaving the configuration unchanged. If, however, the Lag 1 item does not belong to the same category as the first target (i.e., if it is a distractor or a masking stimulus), it will trigger an exogenous change in the system's configuration which, as a consequence, will no longer be optimally tuned to the category of stimuli defined by the first target. The ensuing target (i.e., the Lag 2 item), therefore, will not be processed efficiently even if it belongs to the same category as the first target. This will render it vulnerable to masking by trailing items thus leading to a second-target deficit. The deficit will diminish progressively as the lag is increased, until the first target has been fully processed and the central processor has re-established endogenous control over the system's configuration.

According to this view, the AB deficit arises not from depletion of a limited resource but from a temporary loss of endogenous control over the system's configuration. We refer to this in short as a *temporary loss of control* (TLC). In a sense, such a loss of control can itself be regarded as an instance of resource limitations. But the limitation, in this case, is on the number of tasks that can be performed concurrently rather than on the available amount of depleting attention. What is hypothesized in the resource-depletion account is a reservoir of processing resources that can handle no more than one or, at best, two items before depleting. In contrast, according to the TLC hypothesis, the system can handle a much larger number of items without loss, provided that they fit the current configuration of the filter. The limit, in this case, is not so much the number of items but the fact that the system can assume only one configuration at any given time.

Thus, the kind of limitation postulated in the TLC hypothesis differs significantly from the conventional resource-limitation view. Perhaps the most telling distinction between the two views, however, is that they generate contrasting predictions, which are examined in Experiment 1.

Experiments 1 and 1a

The basic display in Experiment 1 consisted of an RSVP stream of digits that acted as distractors. Inserted near the end of the distractor stream was a three-item *target string* that contained either three letters or two letters separated by a digit, as illustrated in Fig. 1. We refer to these as the Uniform and the Varied string respectively. Observers were required to report all three letters in the Uniform string, and the two letters in the Varied string. Comparing the display sequence in the present study with that in conventional AB studies, the first item in the present target strings corresponds to the first target in AB studies, the second item to the Lag 1 item, and the third item to the Lag 2 item.

Predictions from the two hypotheses can be illustrated with reference to the letters common to both strings, namely, the first and the third items in the Uniform and the Varied strings. According to the resource-depletion hypothesis, the results should reveal a conventional AB deficit in both strings. Accuracy of identification for the third item in the Uniform string should be lower than accuracy for the first item because the available resources are said to be preempted by the first item to the detriment of ensuing items. This should be equally true for the Varied string. If anything, the third item in the Varied string should fare better than the corresponding item in the Uniform string because fewer resources should be required for processing two letters than three. In brief, an AB deficit is expected in both the Uniform and Varied conditions.

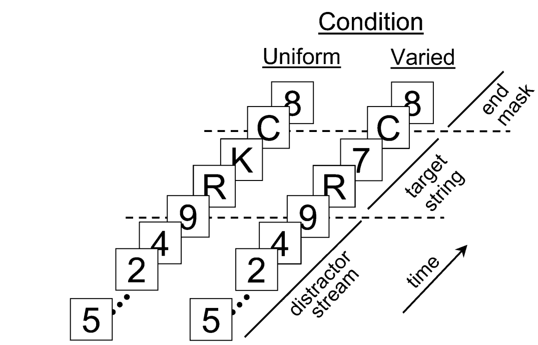

Fig. 1 Schematic representation of the display sequence of stimuli in Experiment 1

According to the TLC hypothesis, on the other hand, an AB deficit should be obtained in the Varied condition but not in the Uniform condition. In the Varied condition, the intervening digit is expected to disrupt the system's configuration in the absence of maintenance signals while the central processor is busy with the first target. This will cause the trailing target to be processed inefficiently, with attendant identification deficit. In contrast, no AB deficit is expected in the Uniform condition because all three targets are letters that match the system's initial configuration and, therefore, do not alter it. In this case, the main limitation is not availability of resources but size of the memory span, which we assume to be at least three letters. The outcome confirmed the TLC hypothesis.

A subsidiary experiment, Experiment 1a, was added to Experiment 1, aimed at checking the generality of the results. In Experiment 1, the distractors were digits and the targets were letters. In Experiment 1a, the categories were reversed: The distractors were letters and the targets were digits.

Experiment 1

Methods

Observers

Eighteen experimentally naïve undergraduate students at the University of British Columbia participated in the experiment. All reported normal or corrected-to-normal vision.

Apparatus and stimuli

Stimuli were displayed on a Tektronix 608 oscilloscope equipped with P15 phosphor at a luminance of 20 cd/m^2. The stimuli consisted of digits and upper-case letters subtending approximately 1° of visual angle at a viewing distance of 57 cm, set by a head rest.

Procedure

At the beginning of each trial, a fixation cross appeared in the center of the screen. Each trial began 500 ms after the observer pressed the spacebar. The display sequence consisted of an initial stream of 5–10 digits selected randomly with replacement from the digits 0–9, with the constraint that the selected digit was not one of the two immediately preceding items. In the Uniform condition, the stream continued with a sequence of three letters chosen randomly without replacement from the English alphabet excepting I, O, Q, and Z, and terminated with a single digit that acted as a mask after the last letter. The display stream in the Varied condition was the same as in the Uniform condition except that the middle letter was replaced with a randomly-chosen digit, as illustrated in Fig. 1. Each item was displayed for 30 ms with an inter-stimulus interval (ISI) of 70 ms. In each condition, observers reported the identity of all letters in the stream, in any order, by pressing the corresponding keys on the keyboard. Therefore, observers identified three letters in the Uniform condition but only two in the Varied condition. The two conditions were presented in separate blocks, each consisting of 100 experimental trials preceded by 10 practice trials. Order of condition was counterbalanced across observers.

Results

Mean percentages of correct responses for the three letters in the Uniform condition were 72.2, 87.2, and 72.9 respectively. The corresponding scores for the two letters in the Varied condition were 85.5 and 59.4 respectively. An analysis of variance (ANOVA) was performed on the accuracy scores for the targets common to both conditions, namely, the first and third items in the target strings. The ANOVA comprised two within-subject factors: Condition (Uniform vs. Varied) and Position in the Target String (First vs. Third). The corresponding scores, averaged over observers, are illustrated in Fig. 2. The analysis revealed a significant effect of Position, $F(1,17) = 13.07$, $p < .01$, $MSE = 172.53$, and a significant interaction effect,

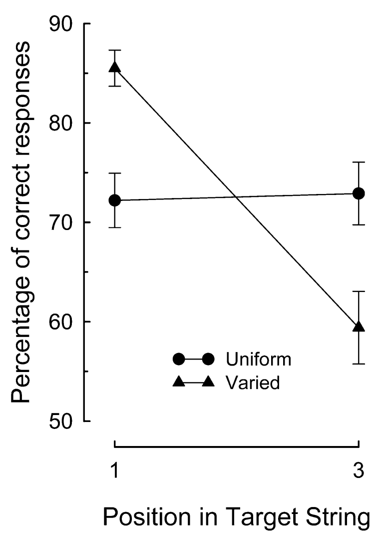

Fig. 2 Mean percentage of correct responses in Experiment 1. In the Uniform condition the target string consisted of three letters. In the Varied condition it consisted of two letters separated by a digit. Observers reported all letters in the strings. The data represent accuracy for the letters common to both conditions, namely, the letters in the first and third positions in the target strings. Error bars represent one standard error of the mean

$F(1,17) = 55.19$, $p < .001$, $MSE = 72.57$. The effect of Condition was not significant, $F(1,17) < 1$.

An AB deficit is evidenced if accuracy for the trailing target is lower than for the leading target. By this measure, the data in Fig. 2 reveal an AB deficit in the Varied condition but not in the Uniform condition. Individual t-tests confirmed that the scores for the leading and trailing targets differed in the Varied Condition, $t(17) = 6.43$, $p < .001$, but not in the Uniform condition, $t(17) = 1.14$, $p > .25$. Indeed, accuracy for the trailing target was lower in the Varied than in the Uniform condition even though the reverse was true for the leading target, possibly reflecting the effect of lighter memory load in the Varied (two targets) than in the Uniform condition (three targets).

Experiment 1a

Methods

The methods in Experiment 1a were the same as in Experiment 1, except for the following. A new group of 18 observers was recruited for Experiment 1a. In addition, digits and letters in Experiment 1 were swapped in Experiment 1a. Thus, in Experiment 1a, the leading distractors, the trailing mask, and the intervening item in the Varied condition were letters, and the targets were digits.

Results

The results of Experiment 1a were very similar to those of Experiment 1, though accuracy was somewhat higher, probably because of a higher a priori probability associated with ten digits than with 22 letters. Mean correct responses for the three digits in the Uniform condition were 87.6, 91.7, and 82.1. The corresponding scores for the two digits in the Varied condition were 91.9 and 74.6. An ANOVA yielded an outcome similar to that in Experiment 1. There was a significant effect of Position, $F(1,17) = 33.61$, $p < .001$, $MSE = 69.46$, and a significant interaction effect, $F(1,17) = 16.86$, $p = .001$, $MSE = 72.57$. The effect of Condition was not significant, $F(1,17) = 1.31$, $p > .25$, $MSE = 35.58$. We conclude that the pattern of results obtained in Experiments 1 is not specific to letter targets and digit distractors.

Discussion of Experiments 1 and 1a

The findings of principal interest in Experiments 1 and 1a were the absence of an AB deficit in the Uniform string, in which the middle item belonged to the same category as the leading item, and the presence of a substantial AB in the Varied string, where the middle item belonged to a different category. Absence of a deficit in the Uniform condition is difficult to reconcile with a resource-depletion hypothesis. The failure is especially telling because the trailing position in the present target strings corresponds to the Lag 2 position in conventional AB studies where the deficit is often largest (e.g., Chun, 1997). Although clearly inconsistent with the resource-depletion hypothesis, which predicts a deficit in both conditions, this pattern of results is consistent with the TLC hypothesis.

Also consistent with the TLC hypothesis is the finding of Raymond et al. (1992) that if the leading target is followed by a blank interval (leaving the target unmasked) and the RSVP stream is then resumed with a varying number of distractors, accuracy of second-target identification is significantly improved (i.e., the AB is reduced). This finding is consistent with the TLC account because, being unmasked, the first target can be processed during the blank interval. As a consequence, the central processor is not busy when the next distractor arrives, and can, therefore, issue the maintenance signals that enable the system to reject distractors and accept targets.

One further aspect of the results in the Uniform condition needs to be considered. In that condition, observers made three responses, the middle of which was more accurate than either the leading or the trailing responses. Thus, accuracy first increased from the leading to the middle response and then decreased to the trailing response. The reason for the increment is unclear. It may be related to the phenomenon known as *residual switch cost* in which an endogenously-set input filter does not reach full efficiency until after it has been engaged by an appropriate exogenous stimulus (Allport, Styles, & Hsieh, 1994; Meiran, 1996; Rogers & Monsell, 1995). Thus, accuracy improved upon the second presentation of a letter. The reasons for the decrement are equally unclear. It may be related to increased memory load with additional targets, or it may be viewed as an incipient, late-developing AB, although this is unlikely because accuracy for the trailing target was within the same range as the leading target (Fig. 2). What is important is that—consistent with predictions from the TLC but not from the resource-depletion hypothesis—accuracy for the trailing target in the Uniform condition in both Experiments 1 and 1a was significantly higher than for the corresponding target in the Varied condition.

Experiment 2

Although the outcomes of Experiments 1 and 1a were consistent with the TLC hypothesis, another interpretation is possible, based on the supposition that the requirement to ignore the middle item in the Varied condition might have contributed to the difficulty in reporting the trailing letter. On this view, the effort involved in suppressing the response to the middle item may have interfered with identification of the trailing item.

We examined this option in Experiment 2, which included a replication of Experiment 1 along with an additional Varied condition in which observers were required to report the intervening digit as well as the two letters. There were three conditions in all, presented in separate blocks, counterbalanced across observers: Uniform (report all three letters), Varied 2 (report only the two letters), and Varied 3 (report the two letters and the intervening digit). Twelve observers participated in the experiment. In all other respects, the methods and procedures were the same as in Experiment 1.

The results, illustrated in Fig. 3, were very similar to those of Experiment 1. Mean percentages of correct responses for the target letters were as follows. Uniform condition: 77.1, 87.9, and 72.1; Varied 2 condition: 86.1 and 55.2; Varied 3 condition: 84.3, 25.5, and 54.8. Clearly, the Varied 2 and Varied 3 conditions yielded virtually identical results, indicating that the AB deficit was not affected by the requirement to ignore the intervening digit. An ANOVA performed on the scores illustrated in Fig. 3 revealed significant effects of Condition, $F(2,22) = 4.60$, $p < .05$, $MSE = 36.75$, Position, $F(1,11) = 23.66$, $p < .001$, $MSE = 362.65$, and a significant interaction effect, $F(2,22) = 25.11$, $p < .001$, $MSE = 50.90$. As was the case in Experiment 1, accuracy in the Uniform condition did not differ significantly between the first and third items, $t(11) = 1.14$, $p > .25$. In brief, an AB deficit was obtained in both Varied conditions but not in the Uniform condition.

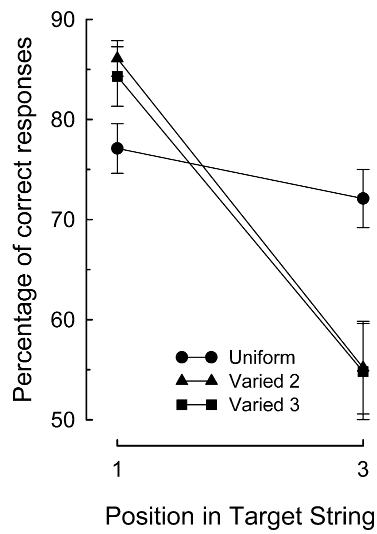

Fig. 3 Mean percentage of correct responses in Experiment 2. The Uniform and Varied 2 conditions corresponded to those in Experiment 1. The Varied 3 condition was the same as Varied 2, except that the observers reported all three items in the target string. The data represent mean accuracy for the letters common to all three conditions, namely, the letters in the first and third positions in the target strings. Error bars represent one standard error of the mean

It is perhaps worth noting that, as in Experiment 1, identification accuracy for the leading item in the target string was higher in the Varied than in the Uniform conditions (Figs. 2, 3). This was possibly due to a lighter memory load in the Varied Conditions. The fact that in the Varied 3 condition observers were required to report all three items in the target string did not seem to add to the memory load because the middle digit was seldom reported correctly (25.5% vs. 84.3% for the first item). It is worth noting, however, that, even though unreported, the middle digit was probably processed to a level sufficient to disrupt the current configuration of the input filter. The important point is that substantial AB deficits were found in both Varied conditions but not in the Uniform condition.

Experiments 3 and 3a

In the preceding experiments, the TLC hypothesis was outlined in broad general terms rather than in fine detail. This was done for two main reasons. First of all, the resource-depletion hypothesis, with which the TLC hypothesis was compared, is itself general and lacking in detail. We have seen how that hypothesis fails to specify even such basic information as the number of items beyond which resources begin to deplete. Secondly, and perhaps more important, the formulation of specific aspects of the TLC hypothesis requires empirical evidence not immediately available in the literature. Experiments 3 and 3a were aimed at providing some of that evidence.

Experiment 3

A question arising from the preceding experiments concerns the type of stimuli that can trigger a reconfiguration of the input filter. For example, in the Varied condition of Experiment 1, the item intervening between the two targets was a digit drawn from the same stimulus population as the leading distractors. We now ask whether the filter reconfiguration evidenced in the preceding experiments depends critically on the intervening item being a distractor, or whether the reconfiguration can be triggered by an item drawn from a different stimulus population.

Experiment 3 was designed to answer this question by replicating Experiment 1 with an additional condition. The new condition, the *Varied (symbol)* condition, was the same as the Varied condition in Experiment 1, except that the item between the two target-letters was not a digit, but one of ten non-alphanumerical keyboard characters similar to those used by Chun and Potter (1995). Thus, in Experiment 3 there were two Varied conditions: The Varied (digit) condition, identical to the Varied condition in Experiment 1, and the Varied (symbol) condition, in which the item between the two target letters was a symbol. To the extent that the impaired identification of the trailing letter in the Varied (digit) condition depends on the intervening item being a distractor, the corresponding impairment in the Varied (symbol) condition should be reduced or eliminated. If, on the other hand, the impairment can be triggered by an intervening item belonging to a category other than the leading distractors, the Varied (symbol) condition should yield much the same results as the Varied (digit) condition.

Method

Observers

Twenty experimentally naïve undergraduate students at the University of British Columbia participated in the experiment. All reported normal or corrected-to-normal vision.

Apparatus and stimuli

Stimuli were displayed as black characters on the white screen of a NEC AccuSync computer monitor operating in a Microsoft Windows environment. The letters and digits were similar to those in Experiment 1. In addition, there were ten non-alphanumerical stimuli: #, $, %, &, (), *, < >, =, ?, and @, whose angular size was approximately the same as that of the digits and letters.

Procedure

Experiment 3 comprised three within-subject conditions: Uniform, Varied (digit), and Varied (symbol). The Uniform and the Varied (digit) conditions were the same as in Experiment 1. The Varied (symbol) condition was the same as the Varied (digit) condition, except that the intervening digit between the two letter targets was replaced by one of the ten non-alphanumerical stimuli, chosen randomly in each trial. The three conditions were presented in separate blocks, each consisting of 100 experimental trials preceded by 10 practice trials. Order of conditions was counterbalanced across observers.

Results and discussion

The results for all three conditions are illustrated in Fig. 4. Mean percentages of correct responses for the three letters in the Uniform condition were 70.9, 86.1, and 73.9 respectively. The corresponding scores for the two letter-targets were 84.5 and 64.4 in the Varied (digit) condition, and 80.0 and 60.5 in the Varied (symbol) condition. An ANOVA was performed on the accuracy scores for the targets common to all three conditions, namely, the first and third items in the target strings. The ANOVA comprised two within-subject factors: Condition, at three levels—Uniform, Varied (digit), and Var-

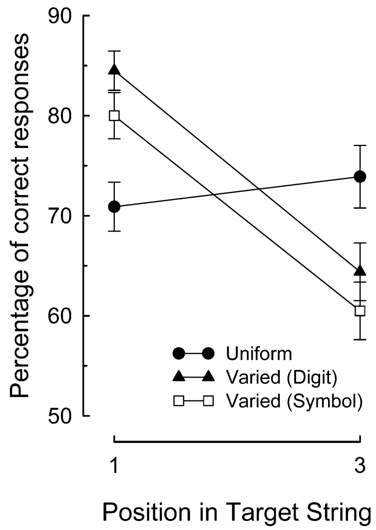

Fig. 4 Mean percentage of correct responses in Experiment 3. The Uniform and Varied (digit) conditions corresponded to the Uniform and Varied conditions in Experiment 1. The Varied (symbol) condition was the same as the Varied (digit) condition, except that the item between the two target letters was a keyboard symbol. Observers reported all items in the target strings. The data represent mean accuracy for the letters common to all three conditions, namely, the letters in the first and third positions in the target strings. Error bars represent one standard error of the mean

ied (symbol); and Position in the target string—First vs. Third. The analysis revealed a marginally significant effect of Condition, $F(2, 38) = 3.20$, $p = .052$, $MSE = 54.43$, a significant effect of Position $F(1, 19) = 15.38$, $p < .001$,$MSE = 288.78$, and a significant interaction effect, $F(2, 38) = 21.24$, $p < .001$, $MSE = 81.82$. Individual comparisons in first-target performance revealed that accuracy in the Uniform condition was significantly lower than that in either the Varied (symbol) condition, $t(76) = 3.49$, $p < .001$, or in the Varied (digit) condition, $t(76) = 5.21$, $p < .001$, which did not differ from each other, $t(76) = 1.72$, $p > .05$. Individual comparisons in third-target performance revealed that accuracy in the Uniform condition was significantly higher than that in either the Varied (digit) condition, $t(76) = 5.13$, $p < .001$, or in the Varied (symbol) condition, $t(76) = 3.66$, $p < .001$, which did not differ from each other,$t(76) = 1.48$, $p > .10$.

The results for the Uniform and Varied conditions in Experiment 3 (Fig. 4) matched those for the corresponding conditions in Experiment 1 (Fig. 2). The finding of principal interest was that the degree of third-target impairment in the Varied (symbol) condition was very similar to that in the Varied (digit) condition. This indicates that reconfiguration of the input filter can be triggered by intervening items drawn from a stimulus population other than that of the leading distractors.

This pattern of results invites a further question regarding the configuration of the input filter: What is the state of the filter when it is reconfigured exogenously by the intervening item? Does the filter change to a new configuration optimally tuned to the class of stimuli exemplified by the intervening item (e.g., digits in Experiment 1)? Or does it simply change into some kind of unfocused state? A series of pilot studies suggested that the new configuration might depend on several factors, notably the structural and semantic similarity between the intervening item and the leading distractors. A systematic examination of these factors is beyond the scope of the present paper. We plan to pursue such an examination in a separate series of experiments.

Experiment 3a

In all the experiments reported thus far, the Uniform and the Varied conditions were presented in separate blocks of trials. This might have allowed for different strategies to be adopted in the two conditions. For example, an attentional gate might have been kept open longer in the Uniform condition, resulting in more accurate identification of the third target. Presenting the two conditions randomly mixed within a single block of trials should disallow such a strategy. On the other hand, whether the two conditions are presented randomly or blocked should be of no consequence according to the TLC hypothesis. This is because loss of attentional control is said to arise not from different strategies but from a mismatch between the leading target and the next item in the Varied condition.

Methods

We examined this issue in Experiment 3a by replicating the Uniform and the Varied conditions of Experiment 1, with the single change that the two conditions were randomly mixed within a single block of trials, instead of being presented in separate blocks. Twelve experimentally naïve observers took part in the experiment.

Results and discussion

Mean percentages of correct responses for the three letters in the Uniform condition were 78.8, 90.6, and 70.6 respectively. The corresponding scores for the two letters in the Varied condition were 86.5 and 58.6 respectively. An analysis of variance (ANOVA) was performed on the accuracy scores for the targets common to both conditions, namely, the first and third items in the target strings. The ANOVA comprised two within-subject factors: Condition (Uniform vs. Varied) and Position in the Target String (First vs. Third). The corresponding scores, averaged over observers, are illustrated in Fig. 5. The analysis revealed a significant effect of Position, $F(1,11) = 33.22, p < .001$, $MSE = 118.13$, and a significant interaction effect, $F(1,11) = 29.47, p < .001$, $MSE = 39.38$. The effect of

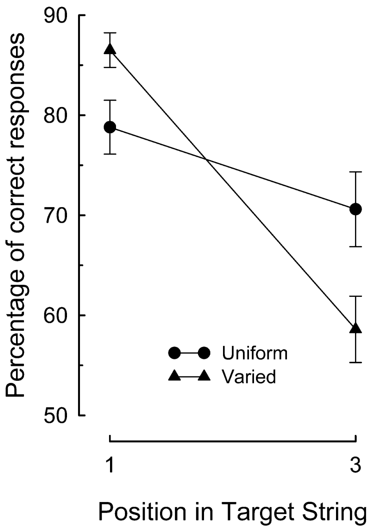

Fig. 5 Mean percentage of correct responses in Experiment 3a. In the Uniform condition the target string consisted of three letters. In the Varied condition it consisted of two letters separated by a digit. Trials in the Varied and the Uniform conditions were presented randomly mixed within a single block instead of separately in independent blocks as in Experiment 1. Observers reported all letters in the strings. The data represent accuracy for the letters common to both conditions, namely, the letters in the first and third positions in the target strings. Error bars represent one standard error of the mean

Condition was not significant, $F(1,11) = 1.35$, $p > .20$. Individual *t*-tests revealed that the leading and the trailing targets differed significantly both in the Uniform condition, $t(11) = 2.39$, $p < .05$, and in the Varied condition, $t(11) = 7.38$, $p < .001$.

The results of the present experiment (Fig. 5) are very similar to those of Experiment 1 (Fig. 2). To examine the effect of blocking vs. mixing of conditions, an ANOVA was performed on the combined results of Experiments 1 and 3a. The ANOVA comprised one between-subject factor (Experiment 1 vs. 3a) and two within-subject factors: Condition (Uniform vs. Varied) and Position in the Target String (First vs. Third). The analysis revealed a significant effect of Position, $F(1,28) = 37.65$, $p < .001$, $MSE = 181.16$, and a significant interaction effect, $F(1,28) = 76.06$, $p < .001$, $MSE = 51.17$. The effect of Condition was not significant, $F(1,28) < 1$. Notably, neither the effect of Experiment, $F(1,28) < 1$, nor any interaction effect involving Experiment was significant: Experiment × Condition, $F(1,28) < 1$; Experiment × Position, $F(1,28) = 1.15$, $p > .29$, $MSE = 181.16$; Experiment × Condition × Position, $F(1,28) = 1.81$, $p > .19$, $MSE = 51.17$.

Clearly, the results of Experiment 3a matched those of Experiment 1 in every important respect. It can be concluded, therefore, that the same pattern of results is obtained whether the Uniform and the Varied conditions are mixed randomly within a single block or are presented independently in separate blocks. This outcome confirms expectations based on the TLC hypothesis.

General discussion

Throughout the present work, an AB deficit was obtained when the stimuli in the target string were drawn from different categories but not when they belonged to a single category. This finding is inconsistent with models based on the allocation of limited resources to the leading target at the expense of ensuing targets. According to these models, stimuli presented after the first target should reveal a deficit arising from depletion of processing resources. In fact, no such deficit occurs when the stimuli belong to a single category.

The present findings, on the other hand, are explained naturally by the TLC account. We propose that the initial processing of incoming stimuli is governed by an input filter that is configured to pass target items and exclude non-target items. The filter is maintained by signals from a central processor that can perform only one function at a time. Prior to the appearance of the first target, that function is to monitor the stream for the presence of a member of the target set of items. As soon as the central processor has detected an item from the target set and has begun to identify it, the maintenance signals for this initial monitoring task are discontinued, and the input configuration comes under exogenous control by the trailing stimuli. If the next stimulus belongs to the same category as the first target, the filter's configuration remains unaltered, the stimulus is processed efficiently, and accuracy of recall is limited only by the short-term memory span.

If, however, the trailing stimulus belongs to a different category from the first target, two things will happen. First of all, the trailing stimulus will take longer to process than a same-category stimulus because it does not match the configuration of the input filter. This delay makes the trailing item vulnerable to masking by the next item in the stream (Giesbrecht & Di Lollo, 1998). Secondly, the configuration of the input filter will be disrupted, so that even ensuing stimuli belonging to the same category as the first target will be processed inefficiently.

An alternative to the TLC account should be considered, based on the possibility that the three-letter string in the Uniform condition in Experiment 1 (or the three-digit string in Experiment 1a) might have been combined in a single perceptual unit through temporal grouping. The resulting "chunking" of items might have eased memory load in the Uniform condition, resulting in more accurate recall than in the Varied condition.

However, before taking the "chunking" hypothesis to be a serious theoretical candidate, several considerations are in order. First of all, the present findings show that for chunking to occur, two conditions had to be met: The items must belong to the same category (letters or digits) and they must be presented within a 300-ms temporal window, that is, both categorical similarity and temporal contiguity were required. Temporal contiguity alone is not sufficient or else chunking would have also occurred in the Varied condition. Furthermore, the literature reveals that the known temporal grouping mechanisms in vision, such as temporal integration, are not effective at temporal intervals beyond about 100 ms (Coltheart, 1980). This means that some as yet unspecified temporal mechanism is at work, grouping items of the same class over these rather long intervals. Furthermore, if such a grouping mechanism was influenced by the same factors as the present study suggests as critical for the maintenance of the input filter, the "chunking" hypothesis would merely turn out to be a redescription of the TLC account.

Secondly, the chunking hypothesis would require further assumptions to explain why the Uniform condition led to higher accuracy for the middle of three items than for the leading and trailing items. In the TLC account this has a plausible explanation in that an endogenously-set input filter (e.g., for letter targets) does not reach its full efficiency until after it has been engaged by an appropriate exogenous stimulus (i.e., the first letter; Allport, Styles, & Hsieh, 1994; Meiran,1996; Rogers & Monsell, 1995). No similar explanation of this aspect of the data emerges from the chunking hypothesis. In brief, several aspects of the chunking account are in need of independent verification before it is ready to be presented as a real competitor to the TLC account.

Besides explaining the present data, the TLC hypothesis provides a coherent account of other major findings in the AB literature. Both the occurrence of Lag 1 sparing and its absence are explained on the principle that Lag 1 sparing occurs only when the second target fits the input configuration set up for the first target. Consistent with this principle, Lag 1 sparing is found when the two targets belong to the same category, but not when they differ from one another in two or more dimensions (Visser et al., 1999).

Also consistent with the present view is the finding that the AB deficit vanishes, or is much reduced, when the distractor item presented directly after the first target, known as the +1 item, is replaced with a blank frame. Removal of the +1 item allows the configuration of the input filter to remain unaltered. Thus, provided that the target in the Lag 2 position belongs to the same category as the first target, it will be processed efficiently, obviating an AB deficit. In a real sense, this is equivalent to deferring Lag 1 sparing to Lag 2.

Recent findings regarding the role of object representations in the AB are also encompassed naturally within this conceptual framework. For example, Pinilla and Valdes-Sosa (in press) found that AB magnitude was greater when the two targets appeared in different objects than when they were part of the same object. Our account of this finding is that the initial configuration of the input filter was set to optimize target detection within one object. If the second target was part of a different object, it did not fit the extant configuration, and accuracy suffered accordingly.

Similar results have been reported by Raymond (2003) who presented an RSVP stream consisting of changing perspective views of the same basic object: A line-drawn trident. The two targets consisted of novel features (a thicker line and an additional line respectively) added to the pattern in the corresponding frames of the trident stream. No AB deficit was found in this condition. In contrast, a substantial AB was obtained when the trident pattern in the first-target frame was replaced with an arrowhead pattern, thus introducing a new object in the stream of tumbling tridents.

Absence of an AB in the tridents-only condition led Raymond to conclude that a second-target deficit is avoidable if the RSVP stream contains the same basic object throughout. By the same token, the presence of an AB in the arrowhead condition was said to be triggered by the requirement to create a new object representation for the arrowhead pattern. Although entirely adequate within Raymond's study, such an object-based scheme fails to account for the absence of an AB in the present Uniform conditions where each letter was a distinct new object requiring the creation of a new and different representation.

Raymond's results are encompassed easily within the TLC framework because its predictions about the AB are based not on object similarity in an observer-independent sense, but on how closely various objects are related to the input filter that the observer has established in order to perform the RSVP task. This aspect of the TLC account is illustrated elegantly in a recent study in which observers were asked to monitor a rapid sequence of words for a target word that defined 'human jobs performed for payment,' such as *banker* or *waitress* (Barnard, Scott, Taylor, May, & Knightly, 2004). Two critical types of foil words could appear prior to the target word in the stream: Words denoting human activities done without payment (e.g., shopper, tourist) and words denoting human artifacts (e.g., television, freezer, park bench). The main finding was that target words were missed more frequently when they were preceded by words denoting unpaid human activities than when they were preceded by words denoting human artifacts. This is consistent with observers setting an input filter to select words denoting human activity, but being unable to tune this filter more finely to select only words denoting human activities for pay.

In summary, the present study demonstrates that resource-limited models cannot account fully for the findings in the AB literature. Instead, the limiting factor appears to be the maintenance of control over the nature of the input filter to the visual system. The system appears to be vulnerable to an exogenously-triggered loss of control, which in a typical AB experiment occurs during the period when the first target is being identified. During this time an item that is not part of the target class easily disrupts target-monitoring mechanisms. It is notable that while the system is monitoring for the presence of this first target, and therefore viewing the leading stream of distractors, this exogenously-triggered loss of control does not occur. The difference between the present TLC account and previous resource-limited accounts is underscored most strongly by the finding that up to three different target items can be processed, without any apparent deficit, provided that these items all satisfy the criteria of the input filter that the participant has established to perform the task at hand.

Acknowledgements This work was sponsored by grants from the Natural Sciences and Engineering Research Council of Canada (VDL and JTE), and by the Japan Society for the Promotion of Science (J-iK). We acknowledge the expert assistance of Mark Rempel in the collection and analysis of the data.

References

Allport, A., Styles, E. A., & Hsieh, S. (1994). Shifting intentional set: Exploring the dynamics of tasks. In C. Umiltà & M. Moscovitch (Eds.), *Attention and performance XV*, (pp. 421–452). Cambridge, MA: MIT/Bradford.

Baddeley, A. (1986). *Working memory*. Oxford: Oxford University Press.

Barnard, P. J., Scott, S., Taylor, J., May, J., & Knightly, W. (2004). Paying attention to meaning. *Psychological Science*, *15*, 179–186.

Broadbent, D. E. (1958). *Perception and communication*. Oxford: Pergamon.

Cherry, E. C. (1953). Some experiments on the recognition of speech, with one and with two ears. *Journal of the Acoustical Society of America*, *25*, 975–979.

Chun, M. M. (1997). Temporal binding errors are redistributed in the attentional blink. *Perception & Psychophysics*, *59*, 1191–1199.

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

Coltheart, M. (1980). Iconic memory and visible persistence. *Perception & Psychophysics*, *27*, 183–228.

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

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

James, W. (1950). *The principles of psychology, Vol. 1*. New York: Dover. (Original work published in 1890).

Jolicoeur, P., & Dell'Acqua, R. (1998). The demonstration of short-term consolidation. *Cognitive Psychology*, *36*, 138–202.

Kawahara, J., Zuvic, S. M., Enns, J. T., & Di Lollo, V. (2003). Task switching mediates the attentional blink even without backward masking. *Perception & Psychophysics*, *65*, 339–351.

Meiran, N. (1996). Reconfiguration of processing mode prior to task performance. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 1423–1442.

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

Pinilla, T., & Valdes-Sosa, M. (in press). Attentional shift time and scene organization: Not all blinks are equal. *Psychological Science*.

Potter, M. C., Chun, M. M., Banks, B. S., & Muckenhoupt, M. (1998). Two attentional deficits in serial target search: The visual attentional blink and an a modal task-switch deficit. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*, 979–992.

Raymond, J. E. (2003). New objects, not new features, trigger the attentional blink. *Psychological Science*, *14*, 54–59.

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.

Rogers, R. D., & Monsell, S. (1995). Costs of a predictable switch between simple cognitive tasks. *Journal of Experimental Psychology: General*, *124*, 207–231.

Shallice, T. (1994). Multiple levels of control processes. In C. Umiltà & M. Moscovitch (Eds.), *Attention and performance XV* (pp. 395–420). Cambridge, MA: MIT.

Shapiro, K. L., Raymond, J. E., & Arnell, K. M. (1994). Attention to visual pattern information produces the attentional blink in RSVP. *Journal of Experimental Psychology: Human Perception and Performance*, *20*, 357–371.

Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing. II. Perceptual learning, automatic attending, and a general theory. *Psychological Review*, *84*, 127–190.

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs*, *74*(11, Whole No. 498).

Treisman, A. (1960). Contextual cues in selective listening. *Quarterly Journal of Experimental Psychology*, *12*, 242–248.

Visser, T. A. W., Bischof, W. F., & Di Lollo, V. (1999). Attentional switching in spatial and non-spatial domains: Evidence from the attentional blink. *Psychological Bulletin*, *125*, 458–469.

Ward, R., Duncan, J., & Shapiro, K. (1996). The slow time-course of visual attention. *Cognitive Psychology*, *30*, 79–109.