Bangor University

DOCTOR OF PHILOSOPHY

An investigation into the relationship between object file continuity and modulation of the attentional blink (AB)

Kellie, Frances Jean

Award date:
2001

Awarding institution:
University of Wales, Bangor

Link to publication

# UNIVERSITY OF WALES

PRIFYSGOL CYMRU

AN INVESTIGATION INTO THE RELATIONSHIP BETWEEN OBJECT FILE CONTINUITY AND MODULATION OF THE ATTENTIONAL BLINK (AB)

By

Frances Jean Kellie

A thesis submitted to The School of Psychology, University of Wales, Bangor, in partial fulfillment of the requirements for the Degree of Doctor of Philosophy.

11th December, 2001

i

Summary

When asked to identify two targets embedded within a rapid consecutive stream of visual stimuli, observers are less able to identify the second target when it is presented within half a second of the first. This divided attention deficit has been termed the attentional blink (AB). Object files are temporary episodic representations thought to provide the basis for object constancy in our perceptual world. Object file research has revealed benefits (faster reaction times) when the same object file can be used across successive visual displays. This contrasts with slower reaction times when there is a mismatch between successive visual displays, possibly because the process of generating new object files is attentionally demanding. In this thesis RSVP methodology was used to investigate the relationship between the AB and object files, and specifically the relationship between the degree of object file continuity inherent in the visual streams and the AB magnitude. The full stream (24consecutive stimuli) experiments revealed an inverse linear relationship between the degree of object file continuity and AB magnitude. High levels of object file continuity were associated with almost complete attenuation of the AB, whilst partial disruption of object file continuity led to less attenuation and complete disruption of continuity led to a robust AB effect. Furthermore, these outcomes were not due to differences in perceptual masking across the stream types. Several of the experiments described here were designed to identify which aspect of the full streams provided the locus of object file continuity and the associated AB modulation. No direct evidence was found to support the hypothesis that items preceding T1 serve this function (Raymond and Sorensen,1994), although such a position cannot be ruled out. Evidence presented here suggests that an important locus of object file continuity in the full stream experiments was due to items intervening T1 and T2; the underlying process at this locus is consistent with Kahneman et al.'s (1992) notion of 'recency' in object file reviewing. This thesis also provides evidence that, when present, distractor items (other than the masks that accompany T1 and T2) are implicated in determining object file continuity and AB modulation.

Acknowledgements & Dedication

For Mum, Dad and Kevin...eek!

Table of Contents

Abstract

This review of the literature begins by considering the nature of attention and specifically divided attention, followed by theories of dual task interference and factors which affect our ability to divide our attention. Some methods for studying divided attention over space and time are then reviewed. Particular emphasis is placed on one temporal method - rapid serial visual presentation (RSVP) - and how it has been used to investigate the deficit associated with dividing one's attention between two targets in close temporal proximity (i.e. the attentional blink, hereafter termed AB). Theoretical accounts of the AB are then discussed along with those experimental manipulations that have succeeded in reducing or eliminating the effect. I move on to consider the role of objects in divided attention and the notion of object files is introduced. This includes object file theory and relevant research which has looked at the nature of object files and what information they might contain. I then look at ways in which the concept of object files can be linked to the AB and consider research which has begun to empirically test such a relationship. In the final section the aims and objectives of this thesis are discussed.

The nature of human visual attention

Dynamic visual information which is projected onto the retinae during natural vision is intrinsically varied in wavelength and luminance across both space and time (Shapiro & Raymond, 1994). From this complex data stream the brain must be able to formulate accurate perceptions and instigate appropriate behavioural responses. To meet these requirements, high level attentional processes along with low level sensory mechanisms work together to control the visual information that reaches the visual cortex. It has been suggested that such regulatory mechanisms may involve heightening (excitation) or reducing (inhibition) visual processing (Shapiro & Raymond, 1994). The construct of *attention* is taken here as a group of neural processes which, when faced with multiple competing perceptual inputs, permit the selection of one particular input over the others (Shapiro, Arnell, & Raymond, 1997).

Divided Attention

Divided attention can be defined as the taking in of multiple sources of sensory information simultaneously (or over a very short period of time) and can be contrasted with selective attention whereby a single source of information is selected from among multiple sources (Pashler, 1998). However, Mordkoff (1999) defines the latter - selecting a single target input from among many processed inputs - as divided attention also.

Accounts of dual task interference

When we attempt to divide our attention between two tasks simultaneously it is often the case that dual task interference occurs where performance in one task (or both) will diminish. There are a number of accounts as to the nature of this interference. *Capacity theories* of dual task interference draw on the notion of a finite pool (or pools) of mental resources. Kahneman (1973) proposed that the amount of mental work that an individual can perform is governed by an overall limit and that dual task (or multiple task) interference only occurs when the total mental effort required exceeds the total capacity of available resources. Other capacity theorists (e.g. Wickens, 1984) also believe that resources are shared between tasks but propose not one, but multiple pools of resources. Pashler (1998) notes that whilst such theories vary in detail, they generally share two key assumptions. The first is that two tasks can be performed in parallel but the amount of mental resources allocated to each task will govern task performance. The second is that individuals can (mostly) exercise control over how these resources are allocated.

In contrast to capacity theories, *bottleneck theories* (also termed single channel theories) adopt a more discrete approach towards dual task interference (Pashler, 1998). According to such views (e.g. Welford, 1952, 1967) some types of mental operations cannot be conducted in parallel and lead to a bottleneck in processing which may occur at various stages such as during memory recall or during the selection or initiation of responses. Dual task theories based on the notion of *crosstalk and similarity* concentrate on the type of processes required in order to

perform the tasks. If the tasks (or the information arising from them) are similar then both tasks will need to utilise the same processors and consequently performing one task will have an adverse effect on performing the other (Pashler, 1998). Navon and Miller (1987) termed this interference *outcome conflict*.

Factors influencing dual task performance

Before moving on to consider empirical methods for studying divided attention it is useful to consider factors which are thought to influence the degree of interference occurring in a dual task situation. These factors (which include practice, instructions, and task similarity) have mostly emerged from studies in which the tasks were continuous in nature (e.g. copy typing or visual-manual tracking), and where investigators monitored aggregate performance over a given period of time (often several minutes) (Pashler, 1998).

In terms of practice, research suggests that highly practised tasks can lead to comparable performance in dual and single task conditions. For example, Spelke, Hirst, and Neisser (1976) found that after extensive training, participants were not only able to take dictation whilst reading unrelated text but that their reading speed and comprehension of the text were unaffected. Where participants do face dual task interference, it appears that performance on each task (the resource allocation) can be modulated according to instructions as to how much emphasis should be placed on each task. It has been shown that this results in a trade-off between increased performance on the emphasised task and reduced performance in the other task

(Gopher, Brickner, & Navon, 1982; Sperling & Melchner, 1978). Often, the degree of interference between two tasks can increase when the respective task requirements are too similar. For example, task similarity could be in terms of a common input or output modality (Allport, Antonis, & Reynolds, 1972; Bornemann, 1942; Wickens, 1984; Wickens, Sandry, & Vidulich, 1983).

Methods for studying divided attention over space and time

Empirical investigations of divided attention necessitate the use of tasks that require participants to process multiple stimuli or multiple properties of a single stimulus. Generally, the bulk of research into divided attention has tended to utilise variants of *report* tasks or *visual search* tasks which operate over space or time (Pashler, 1998). These tasks will now be considered with special emphasis on dual task paradigms operating over the temporal (rather than spatial) domain.

Spatial methods for studying divided attention (report tasks and visual search)

Report tasks typically involve presenting a brief display consisting of a number of items and subsequently asking participants to report as many items as possible. This is the basic task which Sperling (1960) used in his classic research identifying a large capacity, but very short term, form of visual memory (which was later termed iconic memory by Neisser (1967). In his experiments, Sperling (1960) presented

very brief displays consisting of 12 letters (presented in one block) and found that participants could only report about five items (despite claiming to have seen the display for a short while after it had terminated). Sperling believed that far more than five items could access this short term memory trace but that the items decay very rapidly unless they are transferred to a more durable memory store. In subsequent experiments Sperling (1960) used either a high, medium or low audio tone 'cue' presented immediately after the offset of the display to indicate whether to report the top, middle or bottom line of the display. He found that nearly all of the items in a cued row were reported accurately, thus demonstrating that the brief memory trace initially contained all twelve items and the attentional cue initiated the selective transfer of one row of letters to a longer term store whilst the other rows were subject to decay. Styles (1997) notes that whilst Sperling's work was primarily concerned with memory processes, his work is relevant to the field of attention because it provides information about the effectiveness of cues which can (or cannot) direct attention (and therefore influence selection) within complex visual displays.

Visual search tasks involve searching a display for one (or more) targets. This might require participants to report whether the target(s) was present or absent (yes/no) or identifying the target(s) from a number of possible targets (this is termed n-alternative forced choice search task) and then measuring accuracy as a function of *set size* (i.e. the number of items in the display). Pashler (1998) believes that studies using speeded visual search task have yielded some of the most illuminating outcomes in recent years. In such studies, the display remains until a response is made. Rather

than using accuracy as the prime dependent measure, speeded search uses reaction

time as its primary measure. Studies using this method are concerned with the length

of time it takes to detect a target(s) (or their absence) as a function of display *set size*.

Speeded visual search tasks have been extensively used to investigate which

combinations of targets and distractors necessitate serial processing (i.e. where each

of the stimuli are processed sequentially) or parallel processing (i.e. where stimuli can

be processed simultaneously). For example, when a search display contains targets

which are different (e.g. colour, size, orientation, luminance) from homogeneous

distractors, then search efficiency is unaffected by the display set size. This search

is termed *pop-out* search and is believed to be parallel in nature. In contrast, serial

searches are thought to be characterised by a positive linear relationship between set

size and RT (Smith, 1962; Treisman & Gelade, 1980).

Duncan and Humphreys' (1989, 1992) theory of visual search and visual

attention is called *attentional engagement theory*. The central tenet of this theory is

the notion of similarity and perceptual grouping: visual search efficiency will be

dependent on how readily the targets and distractors form independent groups. To

examine this Duncan and Humphreys (1989) conducted a number of experiments in

which they manipulated the homogeneity of the distractors as well as target to

distractor similarity. For example, subjects were asked to search for the letter L in an

upright position presented among letter T distractors that were rotated in either in a

homogenous or heterogeneous direction. It was found that maximal search efficiency

occurs when target to distractor similarity is minimal and distractor to distractor

similarity is maximal. Duncan and Humphreys' work demonstrates that search

efficiency depends of the nature of the task and that a clear distinction between

parallel and serial search is untenable.

Temporal methods for studying divided attention over time.

Although much less research has been directed towards the temporal allocation

of attention, experimental methods have been devised to study this aspect of

cognition (Shapiro & Raymond, 1994). Two temporal methods will be considered

here, the psychological refractory period (PRP) paradigm and rapid serial visual

presentation (RSVP). After briefly outlining the nature and outcome of the PRP

method, RSVP is covered more thoroughly since this paradigm is used extensively in

the work presented in this thesis.

Psychological refractory period (PRP)

One of the oldest and simplest methods for studying dual task interference over

time is the psychological refractory period (PRP) paradigm (Pashler, 1998). Briefly,

this procedure involves making speeded responses to two stimuli whose onsets are

separated by a temporal interval. When individuals are engaged in processing and

responding to the first target, the PRP reflects a period of time during which

processing and responding to a subsequent stimulus is slowed, especially with short

interstimulus intervals (Pashler, 1994; Welford, 1952). This effect is generally robust

and has been obtained using simple (Telford, 1931) or choice (Creamer, 1963) RT

tasks, and across different response modalities such as manual and vocal responses

(Pashler, 1990), or manual and eye-movement responses (Pashler, Carrier, & Hoffman, 1993). Theoretically, there is considerable evidence to suggest that the PRP represents some sort of bottleneck involving central (response/selection) mechanisms needed in order to process both tasks. Until the processing demands of the first task have been met these processes are less able to accommodate the processing demands of the second task (Pashler, 1998).

Rapid serial visual presentation (RSVP)

Rapid serial visual presentation (termed RSVP by Forster (1970)) is a valuable tool for studying the allocation of attention over time and is analogous to spatial visual search tasks but working across time rather than space (Raymond, Shapiro, & Arnell, 1992). Chun and Potter (1995) posit that RSVP is useful for researching cognitive processes within the temporal domain, chiefly because it enables the experimenter to place stringent controls over stimulus form, duration, interval and presentation rate.

Multiple-task RSVP methodology uses stimuli (such as letters, words, pictures, digits or sentences) which are usually presented in one spatial location in a rapid consecutive sequence (normally between 6-20 items/sec.). Within the stream of stimuli one item is designated as the first target (T1) and this is followed by subsequent targets (T2, T3, etc.). Target items will differ in some way (e.g. colour or identity) to all other items in the stream. Participants are required to identify (or detect) T1 and T2 (the dual task condition) or to ignore T1 and just identify (or

detect) T2 (the single task condition). Such tasks allow the experimenter to examine participants' performance in relation to the interval between targets by manipulating the serial position of subsequent targets in relation to T1 (Broadbent and Broadbent, 1987).

Building on the earlier work of Lawrence (1971), Broadbent and Broadbent (1987) used dual task RSVP methodology and found that correct identification of T1 impeded correct identification of T2 if it was presented within 400 ms. However, as the period of time between the onset of T1 and T2 (referred to as stimulus onset asynchrony, or SOA) increased, T2 task performance improved considerably. Broadbent and Broadbent claim that identification of a single target from among non-targets can proceed every 100 ms or so but once a target has been identified a unique process comes into force which results in a reduced ability to process subsequent targets for a short period of time. Also using multiple-task RSVP methodology, Reeves and Sperling (1986) and Weichselgartner and Sperling (1987) reported significant impairments in processing post-target items to an output stage. In both studies, participants were asked to name a target (a boxed or highlighted digit) and the subsequent three digits. Both studies reported that the target, the first subsequent item and items occurring 300-400 ms after the target were usually reported. However, items appearing 100-300 ms after the target often went unreported.

A replication of Weichselgartner and Sperling's (1987) study was carried out by Raymond et al.(1992) who, despite using different method and stimuli (apart from

Experiment 1), still found an attentional deficit in the form of impaired ability to correctly identify letters presented within 180-270 ms of T1 presentation. To investigate whether this phenomenon was in part due to memory-based reporting difficulties, Raymond et al.(1992, exp.2) made the task less complex - participants were required to identify T1 (a white letter) and detect whether T2 (a black X) was present or absent in the post-T1 stimuli stream (positioned randomly in one of eight possible post-T1 serial positions). A graphic example of a stimulus stream used by Raymond et al. is given in Figure 1.

Once again, in the dual task condition a significant post-target response deficit was found, this time, between 180-450 ms after presentation of T1. The single task condition of the same experiment, where participants were required to ignore T1 and merely detect the presence or absence of T2, did not result in similar performance deficits - thus providing evidence that the post-target deficits found in the dual task condition did not reflect a limitation in memory or sensory processes but in attentional processes (Raymond et al., 1992).

The RSVP studies discussed thus far have used different stimuli and response formats. Nevertheless they all report similar costs: presenting a second target in close temporal proximity to the first target has a deleterious effect on participants' ability to carry out the T2 task. This deficit was termed the 'attentional blink' (AB) by Raymond et al.(1992) who drew an analogy with the visual eye blink (which

occurs when eyes are presented with important stimuli and serves to interrupt visual

processing).
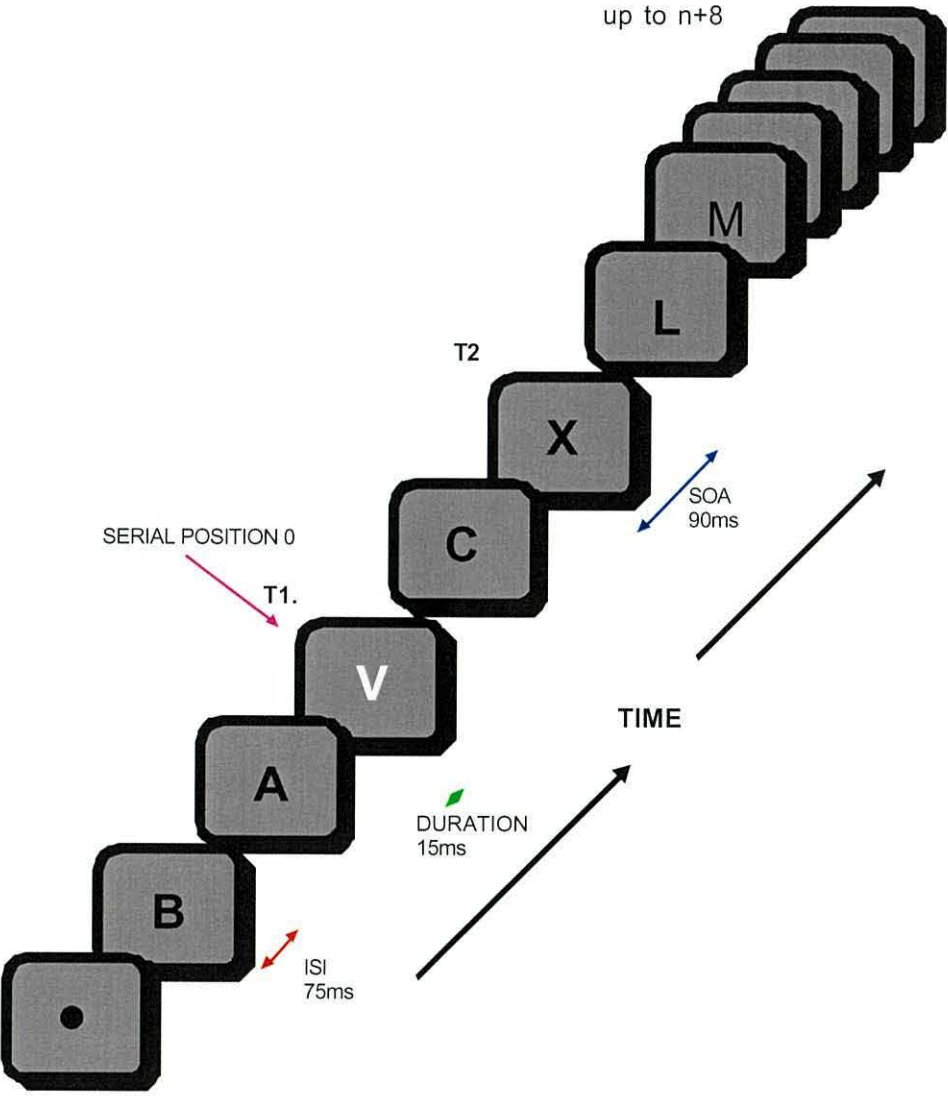
Figure 1.
An example of an RSVP stimulus stream (as used by Raymond et al. 1992, Experiment 2) in which participants were required to report the name of the white letter (T1) and detect whether a black X (T2) was present in the post-T1 stream. (Adapted from Raymond, et al., 1992, p.852).

<u>Illustrating the AB</u>

Usually, when analysing data from an AB RSVP task, participants' ability to correctly perform the second task, having successfully performed the first, is calculated and plotted alongside their performance in the single task condition. A typical outcome revealing the AB effect using this format is shown in Figure 2. An alternative way in which to display the results from an AB experiment is to remove the factor of task type (i.e. single or dual) and use the mean AB magnitude as the dependent variable. An example of this format is shown in Figure 3. This measure is derived by simply subtracting each participant's T2 performance in the dual task condition from that obtained in the single task condition for each of the serial positions of T2. This is a useful strategy to adopt when making comparisons between groups because it takes into account baseline (single task) performance for each individual. Additionally this measure reduces the number of lines displayed on a graph.
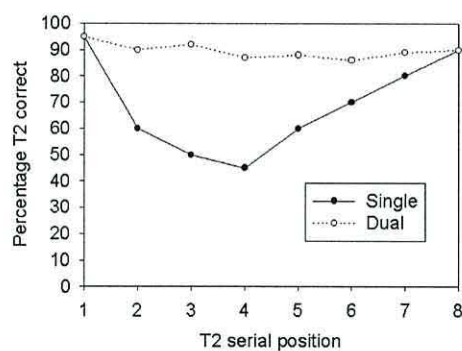
Figure 2
Mean T2 performance in single and dual task conditions plotted as a function of the serial position of T2 after T1. (Adapted from Shapiro and Raymond, 1994, p.161.)

Figure 3.
Mean AB magnitude plotted as a function of the serial position of T2. Negative figures indicate superior performance in the dual task condition compared to single task performance.

In addition to the qualitative pattern of data shown in Figures 2 and 3, the statistical signature of the AB (e.g. within an analysis of variance or ANOVA) is either an interaction between the type of task (i.e. single or dual) and the serial position of T2, if using Figure 2 data format, or a significant effect of the serial position of T2 if one is adopting the derived measure depicted in Figure 3.

Theoretical accounts of the Attentional Blink

In this section the following theoretical models of the AB are outlined: the *interference model* put forward by Shapiro and Raymond (1994) and Shapiro, Raymond and Arnell (1994); the *inhibition model* of Raymond et al.(1992); the *two-stage model* of Chun and Potter (1995); *masking and object substitution* (Enns & Di Lollo, 1997; Giesbrecht & Di Lollo, 1998; Seiffert & Di Lollo, 1997); and Duncan, Ward and Shapiro's (1994) notion of *dwell time*.

Inhibition

Raymond et al.'s (1992) original theoretical explanation of the AB was based on the notion of inhibition. It was suggested that the AB might reflect T1 processing monopolising attentional resources to such a degree that the processing of any T2 attributes is inhibited. This would assist those mechanisms involved in identifying

T1 because features belonging to T2 could not be confused with those belonging to T1, thus preventing T1 conjunction errors. According to the inhibition account, selection occurs early in processing. One prediction that arises from this account is that the AB might be attenuated if the T1 task to be performed was less attentionally demanding. This prediction was tested by Shapiro et al.(1994) in a series of experiments which manipulated task difficulty by using different types of T1 stimuli. Throughout these experiments the T2 task was the same - to report whether a black X was present in the post-T1 stream. Shapiro et al. found, that when T1 was a black 'S' or random white dots a classic AB effect was obtained even though the two T1 target stimuli varied in task difficulty. In contrast, detection of a single gap or distinguishing between a single or double gap in a stimulus stream, which were both very difficult tasks, led to a complete attenuation of the AB. Shapiro et al.(1994) examined the relationship between AB magnitude and T1 task difficulty (as indexed by the d' value for each target's detectability). They concluded that there was no correlation between task difficulty and AB magnitude (although this is a point of controversy which will be considered shortly).

Another prediction one can derive from this account is that no aspects of T2 should be processed whilst limited attentional resources are employed in processing T1. Subsequent work of Shapiro, Driver, Ward and Sorensen (1997) used a three target RSVP paradigm in which T2 was presented inside, and T3 outside, the AB period. It was found that T2 stimuli which could not be identified (because of the AB) were able to prime identification of T3. That is, participants were more likely

to correctly identify T3 if the unidentified T2 was related to T3 than if it was unrelated. This experiment demonstrates that although T2 was not available for report it does achieve high levels of processing (Shapiro, Driver et al., 1997). Converging electrophysiological evidence to this effect comes from an event related potential (ERP) study which measures the brain's electrical activity at the scalp. Luck, Vogel and Shapiro (1996) (see also Vogel,Luck and Shapiro, 1998) presented T2 at the T1+1 (lag 1) position, the T1+3 position (lag3) and the T1+7 position (lag 7). Whilst the T2 responses revealed an AB effect for only the lag3 position, T2 stimuli for all of the lags elicited N400 peak waves which are associated with semantic analysis. Clearly, an inability to report cannot be taken as evidence that high levels of processing are not taking place. These results appear to be problematic for an early selection theoretical position based on inhibition. However, evidence that T2 is processed to high levels despite being unavailable for report is entirely consistent with this theoretical account if one conceptualises inhibition in terms of reduced processing of T2 rather than a complete cessation of T2 processing.

Interference

In 1994, Shapiro et al. adopted an alternative, late selection, theoretical position to explain the AB based on the notion of interference. This account applies aspects of Duncan and Humphreys' (1989) similarity theory of visual spatial search to a temporal model of visual search. According to Shapiro et al. (1994) internal templates for the two targets are matched against perceptual descriptions for stream items. Based on this matching process items are allocated weights within visual short

term memory (VSTM). These weighting resources are limited and decay with time. Since T1 closely matches its template (and because it is the first item to draw on the pool of weighting resources) it is heavily weighted. Similarly, because of the close temporal contiguity of T1's mask (the +1 item) it will also gain access into VSTM and draw from the weighting allocation. However, Raymond, Shapiro and Arnell (1995) suggest that the +1 item enters VSTM not simply because it is contiguous to T1, but that it is also subject to some sort of template matching process that determines whether it enters VSTM, its weighting therein and its consequent ability to interfere with the process of selecting items from VSTM. The greater the similarity between the +1 item and T2, the more heavily the +1 item is weighted. Since T2 closely matches its template, it too will gain entry into VSTM (along with its mask) but the extent to which it is weighted will now depend on the amount of available resource (Shapiro et al., 1994). Thus, at early T2 serial positions, T1 and the +1 item may consume so much of the weighting that T2 is insufficiently weighted, thus rendering it vulnerable to selection errors and yielding an AB effect. At later T2 serial positions, the weighting allocated to T1 and the +1 item will have decayed resulting in the restoration of the weighting resources to the point that T2 can be sufficiently weighted.

Raymond et al. (1995) suggest that the probability that any given item is selected from the VSTM and processed to the level of report is dependent on a number of factors including its weighting, the number of items in VSTM and the degree of similarity between them. There is empirical evidence for a positive linear

relationship between AB magnitude and the number of competing items contained within VSTM. Isaak, Shapiro and Martin (1999) conducted an RSVP study using streams of five items (T1 and its mask, T2 and its mask and the T2+2 item) comprised of random combinations of letters and false font stimuli. On any given trial the number of real letters ranged between two and five. A relationship was revealed between the size of the AB and the number of real letters in the RSVP stream. The work of Isaak et al. shows there to be a relationship between the number of competing items and the size of the AB, but only when the competing stimuli share the same conceptual category. This outcome has been replicated by Taylor and Hamm (1997) using letter or digit T2 stimuli.

Raymond et al. (1995) conducted a series of experiments in which similarity between the stream and the T1+1 item was manipulated. Using traditional letter RSVP streams, they found that the magnitude of the AB was reduced when the +1 item was spatially dissimilar (i.e. displaced from the remainder of the stream) or featurally dissimilar (a random dot pattern). However, McAuliffe and Knowlton (2000) suggest that the random dot pattern may simply not have caused sufficient perceptual interference. In contrast to other categorical manipulations (e.g. Isaak et al., 1999; Taylor and Hamm, 1997), Raymond et al. found no AB attenuation when the +1 item was categorically dissimilar (a digit). Grandison, Ghirardelli and Egeth (1997) note that since Raymond et al. (1995) used letters for T1 and T2, varying the similarity between T1 and the +1 item also varied the similarity between the +1 item and T2. Further investigations into the nature of the interference arising from the +1

item have been conducted by Grandison et al., (1997) and McAuliffe and Knowlton (2000). Since their research impacts on the two dominant theories of the AB (the interference model and the two-stage model) discussion of their work will continue after the two-stage model has been outlined.

Two-stage model

Chun and Potter (1995) attempt to explain the AB in terms of a two-stage model. During the first stage all RSVP stream items proceed in parallel to high levels of analysis (e.g. semantic and featural registration). However, stimuli only become available for report when they have been processed by a second serial processing stage. By this account, the AB occurs because the second stage is still employed in processing T1 and therefore cannot process T2. In support of this position, Chun and Potter report a positive relationship between T1 task difficulty and the size of the AB. They propose that in difficult T1 tasks (e.g. where T1 is difficult to discriminate from the distractor items) T1 occupies the second (serial) stage for longer and therefore there will be an increase in the amount of time during which T2 cannot gain access to this stage. However, something that is problematic for this account is the common finding that the AB deficit does not occur until about 100 ms after T1 (this is termed *lag one sparing*). If, as the two-stage model suggests, T1 ties up the serial processing stage, then it should not be feasible for a T2 item presented at the first serial position to be processed (Pashler, 1998).

<u>The AB and the influence of the T+1 item.</u>

McAuliffe and Knowlton (2000) examined whether the magnitude of the AB effect is determined by featural interference or conceptual interference whilst controlling for stimulus energy (number of pixels). In these two experiments the tasks were to identify a blue letter (T1) and detect the presence of a black Z (T2) among a letter stream. In terms of featural interference the T1+1 item was either a W, a thickened I, or a random pattern of dots (that covered at least the same spatial area as the W). A significant AB was obtained in all conditions but the most featurally complex condition (the W) yielded the largest AB effect. The dot pattern resulted in the smallest AB effect and the thickened I led to an intermediate AB effect. Based on this outcome, McAuliffe and Knowlton suggest that featural complexity, rather than the overall spatial area of the item following T1, is a major determinant of AB magnitude. The role of conceptual interference was then examined by using the letter V as the T1+1 item in either an inverted (the low conceptual interference condition) or an upright position (the high conceptual interference condition). Whilst this manipulation did affect T1 accuracy it did not influence AB magnitude leading the authors to conclude that conceptual interference is not a critical factor in determining AB magnitude, although they do concede that conceptual interference may be observed in paradigms where conceptual characteristics of the stream are more salient (see for example, Chun, Bromberg and Potter (1994) where the targets were defined conceptually).

The interference model (Shapiro et al., 1994) postulates that 4 items enter into VSTM (T1, T2 and their respective masks) having been conceptually identified. Once in the buffer these items compete for retrieval, resulting in a detrimental effect on T2 performance at early SOA's. Since featural not conceptual manipulations of the +1 item affected AB magnitude, McAuliffe and Knowlton's (2000) findings do not support the interference model. Instead the findings support the two-stage model but suggest that the AB occurs because of interference taking place during the first stage where low level features are processed in parallel.

Rather than investigate featural or conceptual interference, Grandison et al. (1997) sought to compare the interference model with the two-stage model in terms of explaining the interference arising from the +1 item. According to the interference account, greater interference occurs when the +1 item is similar to T2 and this is the major determinant of the AB (Raymond et al., 1995). In terms of the two-stage model, Chun and Potter (1995) made the assumption that increased similarity between the +1 item and T1 will lead to increases in the duration of stage two processing for T1 and will yield a larger AB. However, Grandison et al. (1997) claim that the two-stage model fits well with an explanation based on how effective the +1 item is at masking T1. In essence, any +1 item that acts as an effective mask will delay processing of T1 and lead to an AB effect (Grandison et al., 1997). *Masking* is a term whereby target identification is made more difficult due to another stimulus (the mask) occurring in close temporal or spatial proximity.

To test these two models of the AB, Grandison et al. (1997) conducted a series of experiments in which they manipulated the nature of the +1 item and its similarity to T2. For all experiments participants were required to identify a green letter (T1) and detect the presence of a black X (T2) embedded in a black letter stream. When the +1 item was a blank screen (Exp. 4) or a green screen flash (Exp. 5) no AB was obtained. This outcome replicates Raymond et al. (1992) who reported no AB effect when the +1 item was removed (thus removing patterned information from the +1 position) and generally supports the interference model. However, it was found that when the +1 item was a white, high luminance screen flash a robust AB was obtained despite there being no patterned information present. Because there was no patterned information the +1 item should not demand any VSTM resources and - according to the interference model - no AB should occur. Grandison et al. (1997) also demonstrated that an AB effect can be obtained using a +1 item that is visually dissimilar to T1 and T2 - a white rectangle (exp. 1). These outcomes are problematic for the interference account. The two-stage model can explain these outcomes in terms of masking; any +1 item that is an effective mask will increase the amount of time during which T1 monopolises the second stage of processing. By this account, the white screen flash and white rectangle are seen to be adequate masks whereas the green screen flash and blank screen are not (Grandison et al., 1997).

Masking and Object Substitution

Another theoretical explanation of the AB – an adaptation of the two-stage model - is centred on the masks which accompany T1 and T2 and the notion of

object substitution (Enns & Di Lollo, 1997; Giesbrecht & Di Lollo, 1998; Seiffert & Di Lollo, 1997; Brehaut, Enns & Di Lollo, 1999). The cornerstone of object substitution theory is that the nature of visual representations for attended and unattended items are vastly different (Brehaut et al., 1999). Compared to unattended items, items that are attended are more durable and have a greater spatiotemporal resolution (Enns, Brehaut, & Shore, 1996; Moran & Desimone, 1985; Posner, 1980; Rensink, O'Regan & Clark, 1997; Treisman & Gelade, 1980; Tsal, Meiran & Lamy ,1995). Conversely, unattended visual items (such as those appearing during the AB) are more susceptible to substitution by their masks because they are less durable and have lower spatiotemporal resolution (Brehaut et al., 1999).

Integration masking involves simultaneous presentation of the target and mask (i.e. mask and target are superimposed) which are perceived as a single pattern. *Interruption* masking, on the other hand, occurs when processing of a first patterned stimulus is interrupted by the presentation of a second patterned stimulus at the same spatial location (i.e. the mask occurs some time after target presentation) (Brehaut et al., 1999).

Building on the work of Giesbrecht and Di Lollo (1998), Brehaut et al.(1999) investigated the differential effects of integration and interruption masking (using digit masks) on T1 and T2 letter targets in a fully crossed RSVP experiment. In terms of masking T1, both interruption and integration masks were sufficient to produce the traditional U-shaped AB function for T2 performance. However, masking effects

were more specific for T2 where only interruption masking yielded the typical AB function. Integration masking of T2 (irrespective of T1's mask) did not result in an AB. Brehaut et al. argue that these results are problematic for the interference model (Shapiro et al., 1994) which would predict the opposite pattern of results: an integration mask, because of its temporal contiguity to T2, should be more likely to enter VSTM than a delayed (interruption) mask. Brehaut et al. (1999) suggest that the differential effects of T2 masking are accommodated within the two-stage model (Chun and Potter, 1995) in the following way. With integration masking T2 accuracy would be impaired simply through the gradual process of decay, whilst interruption masking could facilitate the preattentive overwriting of T2 by a subsequently presented item (the mask). Therefore one might expect the latter case to result in a more detrimental effect on T2 accuracy because the T2 item may disappear rapidly during this overwriting process (Brehaut et al., 1999).

Whilst proponents of the object substitution model generally support the two-stage model of Chun and Potter (1995) they tend to emphasise the term substitution, in preference to interruption or overwriting, in order to highlight the function of the interruption mask that follows T2 (Enns & Di Lollo, 1997; Giesbrecht & Di Lollo, 1998; Brehaut, et al., 1999). Enns and Di Lollo (1997) suggest that the function of the mask is not merely to terminate target processing. Rather, the mechanisms involved in object recognition seem to turn their focus from processing T2 to processing its mask instead. This claim is supported by false positive reports where

the T2+1 item is reported as being the T2 item (Martin, Isaak, & Shapiro, 1995; Chun, 1997).

Common themes

Whilst the aforementioned theories diverge in terms of the proposed mechanism underlying the AB, Shapiro, Arnell et al.(1997) identify three central tenets which are common to all the theories mentioned here.

•**Tenet 1** - the presence of a T1 mask means that increased attentional resources are required in order for T1 to be processed to the level of report.

•**Tenet 2** - Since attentional resources are limited then the increased capacity required to process T1 means that there are fewer resources available to process T2. Therefore, because T2 cannot attain a durable representation facilitating report it is at risk of decay or object substitution (i.e. by T2's mask). However, this is resolved within approximately 500 ms as the identity of T1 becomes consolidated.

•**Tenet 3** - in distinctly specific conditions, response selection factors will have a further detrimental effect on T2 accuracy resulting from additional pressures imposed on attentional capacity (such as a requirement for speeded responses).

These tenets can be seen as an attempt to bring together diverging theoretical accounts of the AB. Similarly, Isaak et al. (1999) highlight the need for a single theoretical framework that integrates current theoretical models of the AB. Formulation of such a framework has the support of other investigators in the field

such as Kawahara, Di Lollo and Enns (2001)who suggest that the AB deficit represents a multidimensional phenomenon featuring aspects identified in various theoretical models.

Dwell time and the integrated competition hypothesis

The attentional deficit observed in AB tasks has also been conceptualised in terms of attentional *dwell time* during which attention is engaged in consolidating object properties into perceptually whole objects (Duncan et al., 1994). In order to directly measure and plot the time course of interference associated with attending to two objects Duncan et al.(1994) presented just two targets (and their respective masks) separated by a variable SOA (0-900 ms). Regardless of whether the task involved identifying single or double features of two objects a familiar pattern of results was obtained: T1 performance was good but T2 performance was subject to interference for about 300-400 ms. Whilst the notion of dwell time is not strictly a theoretical model, the outcome of the Duncan et al.(1994) study fits well within the framework of Duncan's (1996) integrated competition hypothesis (ITCH) and agrees with other converging neuropsychological and neurophysiological evidence.

According to ITCH attention is not conceptualised as an abstract resource but as, "an integrated state in which multiple brain systems converge to work on the different properties of a selected object and its implications for action" (Duncan, 1996 p.551). The basic premise of ITCH is that incoming visual information leads to the simultaneous activation of neurons in many different sensory and motor brain

systems and results in numerous motor outputs (e.g. speech, eye movements or reaching and grasping). Duncan (1996) notes the selective nature of vision: at any time, despite being bombarded with visual information, only a small proportion actually reaches conscious awareness to the point that it can be identified or used to control behaviour. Clearly, we selectively attend to some things but not others and in attempting to explain the nature of this selectivity ITCH makes three general proposals. First, competition occurs at the neural level within a diverse range of sensory and motor brain areas (both cortical and sub-cortical). Secondly, during this competition increased representation (neural activity) for one particular object will bring about reduced representation for other objects. Within a system, as soon as one object becomes superior over the others (in terms of activation) this superiority is transmitted and adopted by the other brain systems. When this 'integration' between systems occurs, a state of 'attention' is said to exist whereby disparate cortical and sub-cortical systems unite such that they simultaneously operate on the same dominant object. Finally, proposal three relates to the way in which the neural competition is controlled. In order to give an advantage to those objects which are currently relevant, target selective units are primed locally.

Neurophysiological evidence for this competition and priming (for competitive advantage) comes from a study in which monkeys were trained in a visual search task using complex objects (e.g. a triangle and a square) (Chelazzi, Miller, Duncan, & Desimone, 1993). In this study, monkeys were shown a single object cue display which was followed by a blank display and then a two object choice display. During

the blank display, single cell recordings were taken from neurons in the monkey cortex which preferentially responded to one object (i.e. the cell's effective stimulus) rather than the other object (i.e. the cell's ineffective stimulus). Increased cell activation was evident when the cell's effective stimulus was cued to be the target in the subsequent display, but activation was suppressed when the cell's effective stimulus was the non-target object.

A note on the AB and T1 task difficulty

Different theoretical accounts of the AB present conflicting evidence regarding T1 task difficulty and whether it modulates the size of the AB effect. Whilst bottleneck accounts (e.g. the two-stage model and object substitution) believe that T1 task difficulty does affect the size of the AB effect, interference theory believes that it does not. Evidence can be found to support either position (McLaughlin, Shore and Klein, 2001). McLaughlin et al. attempted to clarify this argument by using a T1 difficulty manipulation that affected the perceptual quality of the target (making it more or less difficult to encode). They used displays consisting of T1 and T2, (black letters) along with their respective masks (pattern masks made from parts of digits). Three levels of target difficulty (easy, medium and hard) were manipulated from trial to trial by varying the duration of the T1 and its mask such that the total duration of T1, the ISI (fixed at 15 ms) and T1's mask was always 105 ms. The T1 difficulty manipulation did affect T1 performance but did not influence the AB. In light of this, McLaughlin et al. suggest that AB magnitude might only be affected by difficulty manipulations influencing central (post-perceptual) processing resource allocation,

such as those involving task switching (Allport, Styles, & Hseih, 1994; Potter, Chun, Banks & Muckenhoupt, 1998; Rogers & Monsell, 1995) or response selection (Jolicoeur, 1998).

The AB - a robust effect

Many empirical studies have used multiple-task RSVP methodology, and a diverse range of stimuli and response formats, to investigate the nature of the AB. Generally it has proven to be a robust phenomenon (Shapiro & Raymond, 1994). However, some experimental manipulations have led to complete attenuation of the AB effect. For example, Shapiro, Caldwell and Sorensen (1997) found that when T2 was the participant's own name the AB effect was eliminated.

Two other AB attenuating experimental manipulations, very pertinent to this research, can be found in a series of experiments which were principally interested in the relationship between AB magnitude and T1 task difficulty (Shapiro et al.,1994). It was found that when T1 was detection of a single gap or distinguishing between a single or double gap in a stimulus stream, complete attenuation of the AB was observed. Apart from concluding that there is no correlation between task difficulty and AB magnitude (although see previous section on the AB and task difficulty) Shapiro et al. (1994) noted that attention must be given to objects or patterned information for the AB to occur.

A recent study by Sheppard, Duncan, Shapiro and Hillstrom (in press) found that attention to patterned information is neither necessary nor sufficient to obtain an AB. When participants were required to judge the duration of a letter T1 no AB was obtained. Sheppard et al. also found no AB when T1 was to distinguish between a single (duration 80 ms) or double gap (160 ms) in the stimulus stream, thus replicating the results of Shapiro et al., 1994. However, the AB was reinstated when T1 was to distinguish between gaps of a much longer duration (440 ms and 880 ms). Sheppard et al. suggest that in order to obtain an AB the necessary factor is the deployment of attention to a new perceptual event. They argue that when gaps in the RSVP stream are short enough to yield the percept of a single stream that briefly stalls at the gap (T1), the AB is attenuated but AB attenuation is no longer observed when the gap duration is long enough that the perception of a single stream is lost and attention must be directed toward the pause that separates the two streams.

Objects are special in visual perception

Empirical evidence supports the proposition that objects serve an important function within the perception of visual scenes, (Gordon & Irwin, 1996). For example, Duncan (1984) presented his participants with either a large or small rectangle with a gap on either the right or left side. Superimposed onto the rectangle was a dotted or dashed line which could be slanted towards the left or right.

Participants were required to make two judgements about the same object (e.g. large rectangle, gap on right) or to make two judgements about different objects (e.g. gap on right, dashed line). Duncan (1984) found that when participants were asked to make judgements about different objects, they were less accurate. This outcome can be explained in terms of Duncan's (1996) integration competition hypothesis (ITCH): In the case of one object, there are no competing inputs and multiple brain systems can easily converge and operate on the single object. In the two object condition, the reduced accuracy might reflect increased attentional dwell time whereby multiple brain systems converge to coalesce the properties of each object. ITCH does allow for more than one object to be processed simultaneously but this may come at a cost (interference), especially if both objects require common processing (Duncan, 1996).

Baylis and Driver (1993) also found benefits of one object judgements over those involving two objects. Their study used an identical visual display for one and two object cases. To achieve this they used ambiguous figures and used colour to manipulate the participants' perceptual set such that the critical edges involved in the task were assigned to either one or two objects. The task in this experiment was to judge which of two apices were lower in the display. Despite the fact that the displays were identical, Bayliss and Driver (1993) found costs (in terms of longer RT's) where the displays were perceptually parsed into two objects rather than one, and they suggested that this two object cost must result, not from low level (bottom up) stimulus characteristics such as spatial frequency but reflect top-down attentional factors.

In another study, Watson and Kramer (1999) asked observers to detect the presence of two features that could either both appear on one object or be presented on two separate objects. The target features were a bend or an opening that could appear at either end of two objects that resembled spanners. Despite equating the spatial separation for both of the target features (when present) across same and different object conditions, the familiar one object advantage was observed.

Whilst one object advantage appears to be a robust finding, Davis, Driver, Pavani and Shepherd (2000) claim that this reflects the amount of information being processed rather than a limitation in the number of objects that can be processed simultaneously. It has been suggested that attending to one part of an object automatically spreads attention to other, task irrelevant, areas of the same object (Egly, Driver, & Rafal, 1994). In light of this, and since previous studies employed two objects of a similar size and shape, the two object condition necessitated attending to a visual area twice as large and complex as the one object condition (Davis et al., 2000). In order to test this assumption Davis et al. (2000) used displays consisting of one large object and two smaller objects to demonstrate that the single object advantage disappears if both conditions have the same surface area. When factors such as surface area and complexity are equated between conditions, two objects can be attended to as easily as one.

If objects play an important role in visual perception then there is also a necessity for a mechanism which serves to maintain object information over time (Gordon & Irwin, 1996). The need for such a mechanism is illustrated in perceptual constancy: when we make eye movements or when objects move, despite the resulting changes in the size, shape and position of the retinal image, we retain a sense of perceptual continuity rather than re-perceive the objects as novel ones (Gordon & Irwin, 1996). Whilst the true nature of such a mechanism is not known, one possible candidate is the notion of the *object file* proposed by Kahneman, Treisman and Gibbs (1992).

Object File theory

Object files are hypothetical constructs which represent a late stage in Treisman's (Treisman, 1986a; Treisman, 1986b) feature integration theory[1] (FIT) which attempts to explain how we perceive whole objects (Styles, 1997). FIT holds that visual features (e.g. colour or orientation) are processed through 3-dimensional stacks of specialised feature maps as follows: activation within each feature map is fed down onto a master map of locations. If an object is categorically distinct then it can be rapidly identified preattentively. However, if object identity depends on a conjunction of features then focused attention is required in order to 'glue' disparate feature map activation. To achieve this, an attentional *spotlight* that automatically

accesses all feature map activation for one location homes in on discontinuous areas of activation within the master map. Finally, this information is fed into a temporary object recognition store (object file) which communicates with long term recognition networks in order to permit identification.

Object files are episodic representations of objects in the real world (Kahneman et al., 1992). Being temporary in nature, object files (also referred to as tokens) are thought to be distinct from more durable, long-term (type) representations used to classify and identify objects. Object files have also been implicated in the accumulation and synthesis of visual information across saccadic eye movements (Henderson & Anes, 1994; Irwin, 1992; Irwin & Andrews, 1996). Theory suggests that an object file is created when we direct our attention to an object in the visual field. This file then collects and maintains information relating to that object and might possibly include all of its defining information such as identity (if it is known) and meaning. The object's physical location, however, is not included within the object file but is instead used as a file index (Kahneman et al., 1992).

The processes of correspondence, reviewing and impletion

Object file theory attempts to explain object continuity through a three-stage process of *correspondence, reviewing* and *impletion.* The process of correspondence establishes whether successive displays involve either the same object in a new location or the appearance of a new object, and this decision is based on low

---

[1] This model has been challenged and modified since its inception but for illustrative purposes the discussion here will be restricted to its more general characteristics.

frequency spatiotemporal information. If information across successive displays

does not correspond then a percept of a novel object in a new location is obtained.

However, if successive displays do correspond then one will perceive the object as

having moved to a new location (Ullman, 1979). When the correspondence process is

complete object file reviewing operations commence. This involves accessing an

active object file, retrieving the information about object attributes and establishing

whether it matches the object currently in view. Object continuity persists when the

information matches. In contrast, when information is incongruent across successive

displays this necessitates modifications to be made to the active object file, or for it

to be discarded and a new object file generated. In both cases object identification

will be slowed (compared to where the same object file can be used) because update-

file or open-file operations are time consuming, presumably because of demands

placed on attentional resources (Kahneman, et al., 1992). The final operation,

impletion, utilises current and reviewed information to create a change or motion

percept which links the successive views (Shepard, 1984).

The nature of object files

The reviewing paradigm

In order to study object file effects and, specifically, to investigate the

integration of information relating to objects which change or move over time,

Kahneman et al.(1992) conducted a series of experiments involving both static objects

and objects in real (or apparent) motion, using a new experimental method which they termed the *reviewing paradigm*. Generally this paradigm involves two visual displays involving two or more different letters. In the first display (the preview field) participants simply viewed letters (in some experiments these were presented either in frames or near vertical bar markers). In the final display (the target display) a single letter was presented (along with frames or bar markers if used). In each of the experiments the target item was selectively linked to an item in the preview field using either location, apparent motion, or real motion. In Kahneman et al.'s stationary displays (Studies 1 and 2) the preview frames remained in the same positions across displays and object file linking depended on the frame in which the letters presented in the preview display.

In Study 3, Kahenman et al. used apparent motion in order to selectively link the target item to the preview display. The preview display consisted of two letters presented one either side of the centre of the screen, and two vertical bar markers positioned above and below the centre of the screen. In the second display the pair of letters were globally displaced to the left or right in order to yield the illusion of the both letters moving left or right (with the direction of motion determined by the relative position of the peripheral visual item in the second display, as reported by Ternus in 1938). Real motion was used in the remaining Kahneman et al. studies (Studies 4-7), whereby preview letters were presented in frames and during a blank linking display the empty frames were seen to move to new positions. All of these experiments shared a common task - participants were required to merely name the

target display letter as quickly as possible and their responses were detected by a

voice activated mechanism which also measured the time between letter onset and

naming (response latency).

Experimental conditions and the calculation of preview benefits

Before commenting on the outcome of Kahneman et al.'s (1992) work is it

essential to consider the nature of the experimental conditions and the way in which

the data were manipulated in order to calculate preview benefits that were presumed

to be, or not to be, object specific. The letter that appeared in the target display was

determined by three experimental conditions (common to all six studies). In the

'Same Object' (SO) condition the target letter was that belonging to the same

perceptual object (or indeed was the same object) presented in the preview field. In

the 'Different Object' (DO) condition the target letter was that belonging to a

different object (or being a different object) appearing in the preview field. Target

letters in the 'No-match' (NM) condition were those that did not feature at all in the

preview field. Figure 4 is a graphical example of the reviewing paradigm used in

Study 4 which is representative of Kahneman et al.'s work.

Figure 4 - Example display of previewing paradigm used by Kahneman et
al.(1992, exp. 4). Note that the target field shows same object (SO)
condition, target would be 'V' in the different object (DO) condition and a
letter other than V or S in the no-match (NM) condition (Adapted from
Kahneman et al., 1992, p.183).

Kahneman et al. (1992) calculated two different types of preview effects. The *non-specific* preview effect represented the benefit (or cost) of being shown the target letter in the preview field and was derived from the difference in RT between the DO and NM conditions. The logic was that RT's in the DO and SO conditions should be faster if the target letter's long-term representation is accessed when that letter appears in the preview field. The second preview benefit was *object-specific* and was derived from the difference in RT between the DO and SO conditions. This represents the benefit (or cost) of the target letter appearing in the same frame in both displays. Here one can see parallels with the single object benefits and different object costs reported by Duncan (1984). Whilst non-specific effects are thought to be due to 'type' priming (priming of long term representations of object features), object-specific effects are thought to reflect individual episodic feature bindings for each of the objects being perceived (i.e. object files) (Hommel, 1998).

Significant object-specific effects were consistently found.

To recapitulate, object file theory suggests that when successive visual displays can be linked to the same object file, as in the SO condition, then object recognition will be faster. Significant object-specific benefits were consistently obtained by Kahneman et al. (1992, studies 1-7) irrespective of the way in which the target and preview displays were selectively linked (static, real or apparent motion) or the experimental parameters (e.g. number of display items, preview duration or interstimulus interval).

Object specific effect is not due to node priming or common location

Kahneman et al. (1992) argue against an explanation of object specificity in terms of activation of nodes in semantic memory (which is the standard account of priming). In node priming there would be no difference in RT between the SO and DO targets because each of the preview items (as well as the position at which they appeared) would be equally primed. Since the object-specific preview effect is derived from SO-DO, benefits associated with object specificity should not occur. Consequently it is claimed that object specificity is not due to node priming (Kahneman et al., 1992). The fact that object-specific benefits were also obtained using displays in apparent motion or real motion, is also problematic for a specificity explanation based on shared location (Kahneman et al., 1992, studies 3-7).

Recency in object file reviewing

Object file theory suggests that when there are strong links between objects in successive displays, attention to an object in the latter display will prompt a review of the object file corresponding to it in the former (most recent) display. Kahneman et al. (1992) presented participants with two successive two-item preview displays and target conditions were again SO, DO or NM but this time in relation to either the first or second preview field. A significant object-specific effect was obtained for the second (most recent) preview display but no preview benefits were apparent for the first preview display. Since earlier experiments showed that specificity is unaffected by long intervening intervals (700 and 590 ms) and, given that there was only a 500 ms interval between the first and second preview fields, an explanation of this result

in terms of iconic memory information for the first preview field decaying can be rejected. The alternative explanation is that information for any frame in the first preview field is simply updated (overwritten) by the letter shown in the corresponding frame in the second preview field. Consequently, the reviewing process (which is instigated by the target) only has access to the most recent object information. Kahneman et al. (1992) suggest this explanation is logical if object files mediate perception and modulate response tendencies.

Object file generation takes time

Whilst object-specific benefits were reported in all seven of Kahneman et al.'s (1992) studies, the conditions under which the size of the benefit was reduced are illuminating in terms of the nature of object files. For example, manipulating the preview field duration revealed that shorter preview exposures resulted in greater non-specific benefits. This suggests that there is a minimum amount of time needed to consolidate object specific links within object files. It is possible that the time needed to achieve object specificity translates to the neural attentional state or dwell time (Duncan, 1996) during which multiple brain systems integrate in order to operate on a single object.

There are limits to the number of objects files we can maintain

Kahneman et al. (1992) also found that the magnitude of the object-specific advantage was inversely related to the number of items present in the preview field, and that this was true regardless of whether the items constituted different letter

types or repeated letter tokens. For example, a two item display yielded a 35 ms object-specific advantage compared to just 6 ms (which was not significant) for a display consisting of eight items. Thus, it appears that separate perceptual representations must be established and maintained for each item in the preview display. As the demand for resources increases the ability to retain object specificity diminishes.

What is in an object file?

Kahneman et al. (1992) claim that object files might contain all of the object's defining information including its identity and meaning. To test this, the reviewing paradigm has been adapted to investigate whether object files contain different types of semantic information (Gordon and Irwin, 1996) and whether object identity information is abstract in nature (Gordon and Irwin, 1996, 2000; Henderson, 1994).

Semantic information

Gordon and Irwin (1996, experiments 4 – 6) investigated whether object files contain different types of semantic information, namely related items, semantic features and categorical information using an experimental paradigm that was only slightly different to that used by Kahneman et al. (1992). Gordon and Irwin used words as stimuli rather than single letters but used the same SO, DO and NM conditions. Here, SO targets were those where the same word appeared in the same

frame in the preview field, DO targets were those where the same word featured within the other frame in the preview field, and NM targets were words which did not feature in the preview field. The majority of Gordon and Irwin's (1996) experiments also included a proportion of non-words. Participants held a microswitch in each hand and were instructed to press one switch if they thought the target was a non-word and the other switch if it was a real word. Reaction times were recorded in each of the conditions and benefits attributable to general priming (i.e. non-specific effects) as well as those associated with object specificity were calculated.

Gordon and Irwin (1996, Exps. 4 and 5) adopted a semantic priming paradigm in order to establish whether object files contain different types of semantic information. To test for related concepts and semantic features a lexical decision task was used in conjunction with different target words. SO targets were related to or synonymous with the same frame item in the preview display. DO targets were words that were related to or synonymous with the word in the other preview frame and NM targets were not related or synonymous with either of the preview words. There was no evidence of benefits associated with reviewing the same object or location.

To investigate whether object files contain categorical information, Gordon and Irwin (1996, exp. 6) presented two preview words which were exemplars from different conceptual categories (e.g. iron and robin) and a single target word that was

the name of a category (e.g. bird). Here the participant's task was to make a speeded decision as to whether either of the preview words was a member of the target category. Again, SO, DO and NM conditions were employed. As with the related concepts and semantic features experiments, no significant object specific benefits were obtained, despite the fact that attending to the preview objects was the crucial task demand here.

In each of the three semantic information experiments conducted by Gordon and Irwin (1996) there were no benefits associated with reviewing the same object or location, despite significant effects of general priming. Based on this, Gordon and Irwin concluded that object files do not contain semantic information - at least that pertaining to related concepts, features or categories. Furthermore, because the reaction times associated with reviewing the same object were consistently found to be longer than in the DO condition (which seems counter-intuitive) Gordon and Irwin suggest that object files might include the integration of some sort of inhibitory component.

Object file representations are abstract

The notion that object files contain abstract information relating to object identity, rather than a description based on physical appearance, was initially investigated by Henderson (1994) who used a preview display consisting of two frames each containing a letter. After a blank linking display (during which the empty frames were viewed) the target display involved the presentation of a single

letter in one frame and a plus sign in the other. In addition to the standard SO, DO, NM conditions, the case (upper or lower) of the target letter relative to that of the preview letters was manipulated. Henderson (1994)[2]. found that the case manipulation had no effect on the size of the observed object-specific preview benefit. Gordon and Irwin (1996, Exp. 3) also compared the object specific benefits derived from preview and target words that were either the same or different case. Significant and comparable object specific benefits were found in both case conditions, once again demonstrating the persistence of object continuity despite changes in stimulus form. Based on this, Gordon and Irwin (1996) suggest that the information which object files maintain about an object's identity is based on the object's abstract physical form.

A recent series of six experiments by Gordon and Irwin (2000) cleverly demonstrate that as long as the identity of the stimulus does not change across displays, object continuity is preserved despite large physical stimulus changes. In one experiment for example, the preview display consisted of two frames each containing a different Snodgrass and Vanderwart (1980) line drawing. The preview frames were presented above and below fixation and during the linking display the empty frames were seen to move to new locations to the left and right of fixation. During the target display a string of letters, which was either a word or non-word,

---

[2]Kahneman et al. (1992) do briefly mention ancillary work in which they changed the case between preview and target displays and the object specific benefit was affected in some conditions but not others. However, Kahneman et al. (1992) failed to report what these conditions were.

was presented within one frame and the participant was required to make a lexical decision regarding the target letter string. When the target was a word it could share the same identity as the picture presented in the same frame (SO condition) or other frame (DO condition) in the preview display, or it could be a different identity to either of the pictures presented in the preview display (NM condition). Gordon and Irwin (2000) found a significant object-specific preview benefit. Whilst this effect did not significantly interact with the position of the preview frame, planned comparisons revealed that this effect was significant for the top, but not the bottom, preview display. This was the case for each of the dynamic displays and the authors suggest that when viewing such displays observers can only maintain a single object file (Gordon and Irwin, 2000).

Response information

There is evidence to suggest that object information need not be solely perceptual but may also include action related information (Hommel, 1998). Whilst Kahneman et al. (1992) did not present evidence that response tendencies were included in object files, they did not rule this out. Hommel (1998) suggests that object files can be subsumed into a far more complex construct – the event file - which contains responses pertaining to previously attended objects.

Gordon and Irwin (2000) reported drastically reduced non-specific preview benefits when the task was a lexical decision compared to when it was simply naming the picture or word. They concede that the larger could be due in part to the

preparation of responses. However, in their experiment involving picture preview displays and a lexical decision target display, response preparation cannot account for the non-specific preview benefit because the nature of the task precludes response pre-preparation (Gordon and Irwin, 2000).

Objects files and the attentional blink

Earlier, I noted how in the AB there are costs associated with dividing one's attention between two targets in a temporal stream for about 500 ms. Similar costs also feature within object file research in terms of slower object recognition times where successive visual displays do not share a common object file. It might therefore be possible to attenuate the AB effect if T1 and T2 share a common object file. To phrase this another way, could the attentional costs observed in the AB be eliminated if attention is divided within an object (or object file) rather than between objects (or object files)? A recent study by Raymond and Sorensen (1994) suggested that this may be the case.

Raymond and Sorensen's (1994) 'trident' experiment

Raymond and Sorensen (1994, exp. 1) used an RSVP stream (with similar parameters to those used in a conventional letter RSVP stream) consisting of a stream of distractor stimuli (tridents) presented in varying orientations to induce the percept of a single rotating trident. The T1 target was a thick bar which could occur on a

trident stimuli (same-object condition), on an arrow (different-object condition), or

not appear at all. Participants were required to report the absence or presence of T1

and, if T1 was present, report whether it was on a trident or on an arrow. The T2

target (which was present on half of the trials) was a trident with a short line attached

to its base. Participants were required to report whether T2 was present or absent.

A graphical illustration of this experiment is shown in figure 5.
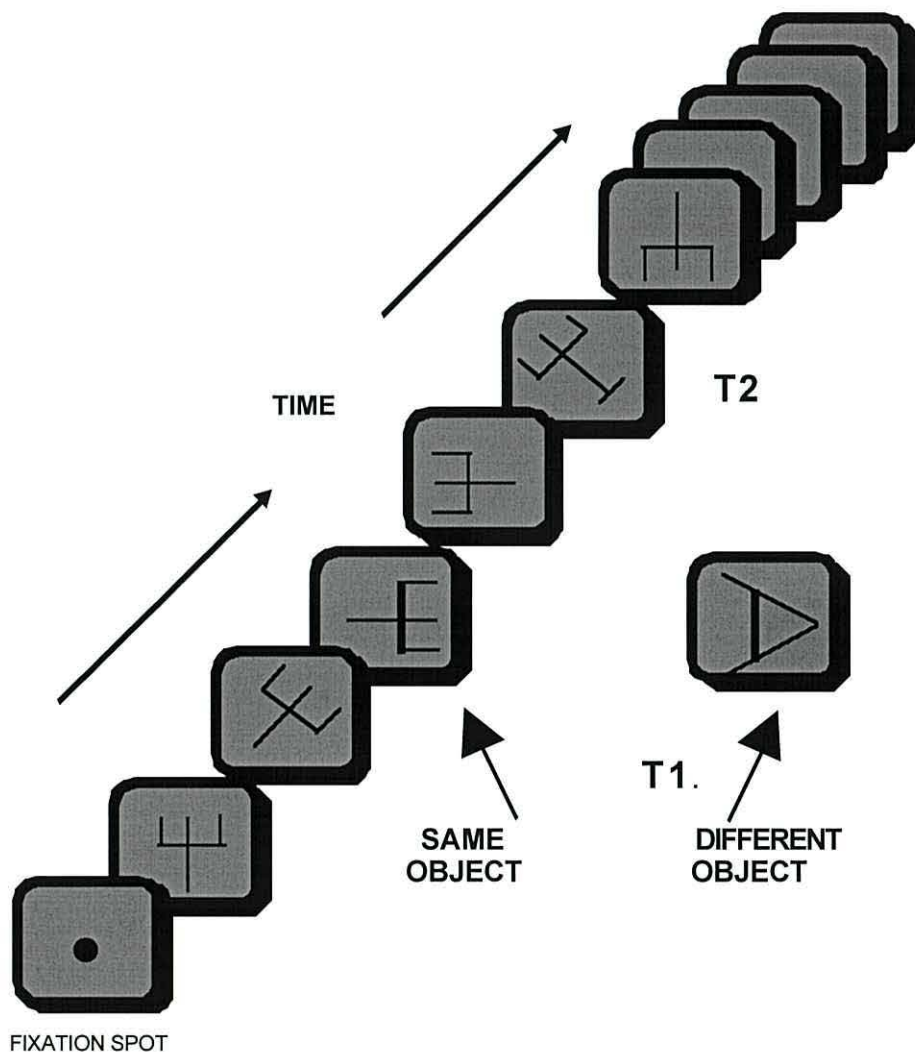
Figure 5.
An example of an RSVP stimulus stream (as used by Raymond and Sorensen, 1994, experiment 1). Here, the T1 task was to report the presence or absence of T1 and, if T1 was present, whether it was a trident or an arrow. The T2 task was to report whether the a trident with an added feature (a foot) occurred in the post T1 stream (Adapted from Raymond and Sorensen, 1994).

In the same-object condition, in which T1 and T2 were features of a single object (the trident) the AB effect was completely attenuated. To explain this, Raymond and Sorensen (1994) suggested that the object file which was initially established when the first distractor target was presented, could be easily updated and utilised for T1 (because it too was a trident). In contrast, in the different-object condition where T1 was a new object (an arrow) a robust AB effect was observed. Here, because the arrow stimulus was a new object, the existing object file could not be easily updated. Instead, the file would need to be discarded and a new object file opened. Since this process is attentionally demanding, the update-file process (which would be required to perform the T2 detection task) may be less efficient until such time as the open-file operation is complete (Raymond and Sorensen, 1994). This scenario is consistent with the observed AB effect in the different-object condition.

Raymond and Sorensen (1994, Exp. 2) went on to investigate whether this outcome was indeed due to the detection of a new object or the consequence of using matching T1 and T2 items. This experiment was essentially the same as the first except that T2 was always the arrow stimulus with a small line placed perpendicular to the arrow tip. It was found that when T1 and T2 were mismatched, but T1 was the same object as the distractors (i.e. the trident), no AB effect occurred. This was taken to show that rapid detection of a feature occurs when it pertains to the object most recently in view. This notion is consistent with Kahneman et al.'s (1992) claim that the most recent object information is reviewed. In contrast, when the two

targets matched, but the T1 item involved detecting a new object (i.e. the arrow) a small attentional blink was obtained for the first T2 serial position. Raymond and Sorensen claim that this outcome reflects a processing deficit bought about because T1 involved the detection of a new object.

Raymond and Sorensen (1994) interpreted their results in terms of object file operations in the following way. Merely updating an object file (i.e. when T1 was a trident) occurs very quickly (under 90 ms) and will not impact upon any subsequent object file relations (e.g. modifying an existing file or opening a new one). However, when open-file operations are utilised (i.e. when T1 was an arrow) there is a refractory period of not less than 270 ms. This refractory period has a differential effect on subsequent object file operations. Update-file operations are less efficient (but only for about 180 ms) whereas open-file operations may be affected for a much longer period.

Aims and objectives of the present thesis

This thesis attempts to further examine the relationship between the AB and object files using RSVP methodology.

- **Chapter 2** preserves object file continuity across streams involving object change (rather than object motion) and demonstrates that object file

continuity leads to AB attenuation. This chapter also demonstrates that disruption of object file continuity leads to an AB effect.

- **Chapter 3** investigates the role of the items intervening T1 and T2 as well as the items preceding T1 and considers whether distractor items – thought to provide object file continuity – are essential for the preservation of object file continuity.

- In **Chapter 4** similar questions are considered but utilising different masks (i.e. those outlined in Chapter 2).

- **Chapter 5** attempts to maximise the efficiency of visual perceptual masking whilst controlling the nature of the targets and masks across each stream type.

- In the final chapter, **Chapter 6**, the results of the experiments presented here will be synthesised and the implications discussed in terms of object file theory. The limitations of this research and avenues for future investigations in this domain will also be considered.

Abstract

Chapter two is comprised of two experiments. Both involve manipulating RSVP streams such that, theoretically, targets can or cannot share a common object file, thus preserving or disrupting object file continuity. The first experiment used a morph movie where a smoking pipe was seen to progressively change into a saucepan and specific movie frames were modified by altering their texture in order to create target stimuli. When the movie was presented sequentially in RSVP (the morph condition) object file continuity was preserved and the AB was attenuated. However, when the same stimuli were presented in a random sequence (the scrambled morph condition), object file continuity was disrupted and a significant AB effect was observed. This effect was relatively short-lived (recovery occurred by 267 ms) reflecting the updating of object files rather than the generation of new ones. The second experiment attempted to force the generation of new object files by embedding the same targets within a stream of random objects. As predicted a significant AB was observed and recovery did not occur until 467 ms.

The full stream morphing experiments

The studies reported here attempted to attenuate the AB effect through linking targets in an RSVP stream to a common object file. Earlier, I noted how object files serve to maintain information about objects as they move or change over time (Kahneman et al., 1992). Whilst Raymond and Sorensen (1994) utilised object file information pertaining to an object's motion over time, the object file information maintained in these experiments related to the object changing state over time.

In order to achieve this a short movie was used in which a smoking pipe was seen to progressively change into a saucepan – a process termed *morphing*. If one adopts the same rationale as Raymond and Sorensen (1994), T1 should adopt the currently active object file by virtue of the preceding movie frames. In the morph RSVP stream where one object changes smoothly into another, each successive frame will be so similar (in terms of spatiotemporal attributes) that the existing object file (the one generated at the start of the stream) should be utilised for each subsequent frame, including T1. If so, then attentional resources would not be required to generate a new object file for T1, and therefore T2 processing should not be hindered from T1 interference (which would, arise from attentionally demanding processes such as open-file or modify-file operations). Consequently, it was anticipated that an attenuated AB effect would be obtained in the morph condition.

A scrambled morph RSVP stream is a good control condition because it uses the same stimuli as the smooth morph stream. Theoretically, this condition should yield an AB effect because object file continuity is disrupted. In the scrambled morph (obtained by randomising the presentation order of all the movie frames except the target frames) there is no longer a percept of smooth object change. Here it was hypothesised that object file continuity would be disrupted and an AB would be observed. Another way in which one can disrupt object continuity (possibly to a greater degree) is to embed the same pipepan targets within an RSVP stream of random objects. In this chapter, the effect of the smooth and scrambled morph RSVP streams was examined in Experiment 1 whereas the effect of using a random object RSVP stream was investigated in Experiment 2.

General Methods

Design

Factorial designs were used for this series of experiments. The dependent variable was RSVP task performance expressed as AB magnitude. The percentage of T2 correct identifications for each serial position of T2 where T1 was correctly identified was calculated for the dual task condition. In the single task condition the number of correct T2 identifications was expressed as a percentage of the total number of single task trials for each serial position of T2. To obtain AB magnitude

T2 performance in the dual task condition was subtracted from that observed in the single task condition for each participant at each SOA.

Participants

All participants were non-dyslexic adults claiming to have normal (or corrected) vision. Participants were recruited from the student participant panel and the community participant database. Community volunteers were paid £5 for a 60 minute session and £7 for a 90 minute session, whereas student volunteers received course credit in return for their participation.

Apparatus

Stimuli were presented using RSVP custom software on an Apple Macintosh 8600/200 power computer. The custom software was authored by Anne Hillstrom, formerly of the University of Wales, Bangor. Stimuli were viewed on an Apple 28 cm RGB monitor using binocular vision from a viewing distance of 60 cm. Illumination in the room was reduced during testing. Participants entered their responses into the computer via the keyboard.

Stimuli & procedure

In order to generate RSVP streams, a morph movie of a smoking pipe changing into a saucepan was created using Morph version 1.1 software. The movie was then dissected into 24 frames using Movie-to-Pict version 2. Specific frames were imported into Adobe Photoshop version 5.0 and texture effects were applied in order

to create the target stimuli. To create T1 images, the twelfth movie frame was modified such that it was comprised entirely of either large dot elements (using pointillist effect, size 6) or small dot elements (pointillist effect, size 3). A different texture effect – patchwork – was used to create T2 stimuli which were comprised of either large squares (patchwork effect, size 4/5) or small squares (patchwork effect, size 1/5). T2 images were adapted from movie frames 13-20. Example target stimuli are shown in Figure 6. Despite changing the textural nature of the elements that made up the targets, the spatiotemporal and colour characteristics of the images from which they were derived were relatively preserved.

Figure 6.
RSVP targets adapted from the morph movie. T1 targets (shown top left and right) were either small or big dots. T2 targets (shown bottom left and right) were either small or big squares.

RSVP streams were displayed at a single location – the centre of a subtended (18.5° by 17.1°) grey field of uniform luminosity. The visual angles subtended by the stimuli varied. Images from the morph movie subtended 2.96° in height and 5.71° in width. The visual angles subtended by the random object distractors (where used) varied in height (from 7.13° to 0.95°) and width (from 6.18° to 2.39°).

In any given trial T2 never preceded T1 and no frame could appear more than once. All participants undertook blocks of single task as well as dual task trials. In single task trials they were required to ignore T1 and only attend to and report T2. In dual task trials participants were required to attend to and report both targets. The order of the blocks was counterbalanced in terms of task type (i.e. single or dual task trials). When commencing the RSVP trials participants were instructed to attend to a centre black fixation dot where the stimulus stream would subsequently be presented. All RSVP trials were self-paced and initiated by the participant depressing the space bar. Participants were instructed to watch each RSVP stream and, once the stream finished, make their keyboard response(s) by pressing keys which had previously been labelled BIG and SMALL. For the dual task trials participants were instructed to make their first response in relation to the first target (i.e. dots) and their second response in relation to the second target (i.e. squares).

Participants were informed that accuracy, and not reaction time, was being measured. Before the experimental trials commenced participants were shown A4 colour printouts of the target stimuli and then three sets of practice trials were

conducted. The first set was comprised of 16 dual task trials (four trials at each T2 serial position 5 – 8, corresponding to SOA's of 333 ms, 400 ms, 467 ms and 534 ms) presented at the slower rate of 133.33 ms per stimulus (8 items per second). The second set consisted of 16 dual task trials presented using the same timings but utilising the earlier T2 serial positions (i.e. four trials at each serial positions 1 – 4, corresponding to SOA's of 67 ms, 133 ms, 200 ms, and 267 ms). The third practice set involved completing one block of 32 dual task trials (four trials at all SOA's) at the faster presentation rate of 66.67 ms per stimulus (15 items per second). Each participant then undertook 320 experimental trials (one set of 160 single task trials and one set of 160 dual task trials). For those experiments utilising seven SOA's rather than eight (all experiments except those in Chapter 2), RSVP streams for the 67 ms SOA were removed from the training and experimental trials. Hence, training set 2 was reduced to 12 trials, training set 3 was reduced to 28 trials and the number of experimental trials was reduced to 240 (one set of 140 single task trials and one set of dual task trials). No blank interstimulus interval was presented between stimuli.

Given that there were two possible T1 responses (big/small dots) and two possible T2 responses (big/small squares), this yielded four possible target combinations for each of the eight T2 serial positions (32 trials). Each combination was repeated five times resulting in sets of 160 trials. This yielded 20 trials at each SOA. Experimental sessions lasted between 60 minutes (for streams without items appearing before T1) and 90 minutes (for full stream experiments).

Participants were replaced if their performance failed to meet criterion. Since 50 percent was chance (i.e. the target could either be big or small) these criteria were set at a mean of 50 percent correct for T1 performance in the dual task trials and for T2 performance in the single task trials. Mean performance refers to performance across all SOA's. Throughout this thesis an alpha level of .05 was used unless stated otherwise, for example where familywise error is accommodated. Other levels of significance (e.g. $< .01$, $< .02$, $< .06$) are occasionally reported for informative purposes only.

Experiment 1 – full stream morph and scrambled morph

Design

This study employed an 8 x 2 mixed factorial design. The within groups factor was the period of time between the onset of T1 and T2 (SOA) having eight levels, 67 ms, 133 ms, 200 ms, 267 ms, 333 ms, 400 ms, 467 ms and 534 ms. The between groups factor was the nature of the RSVP streams (STREAM) having two levels: MORPH where successive stream items gave the percept of a smoking pipe changing into a saucepan, and SCRAMBLED, where the morph movie frames were presented in a random order about T1 and T2.

Participants

Sixteen volunteers were recruited from the student participant panel (5) and the community participant panel (11). Participants were allocated randomly to the morph (5 female, 3 male, mean age 23.6 years, SD = 6.23 years) and scrambled conditions (6 female, 2 male, mean age 24.5 years, SD = 6.74 years).

Stimuli & procedure

An example RSVP stream for each condition is shown in Figure 7. Participants were required to identify the size of the dots for T1 and the size of the squares for T2.

Figure 7.
RSVP streams for the smooth morph (shown on the left) and scrambled morph (shown on the right). Streams were 24 frames long but have been truncated here for practical purposes. T2 is shown at serial position 2 which corresponds to an SOA of 133 ms.

Results and Discussion

The mean percentage of correct T1 responses was significantly higher in the smooth morph condition (97%, SD 5.00%) than in the scrambled morph condition (89%, SD 12.00%), $t(126) = 4.728$, $p < 0.000$ (to gain some measure of how well the computer had randomised the presentation order of the morph movie frames in the scrambled stream two data sets were analysed and the mean and median morph movie frame calculated for specific positions surrounding T1 and T2. These are given in Appendix A).

Mean AB magnitudes were calculated for each SOA and for both stream types. These are displayed in Figure 8 which suggests that the proportion of correct T2 responses was reduced, for the dual relative to the single task, at early SOA's for the scrambled condition but not for the morph condition. That is, there appears to be an AB in the scrambled condition but not in the morph condition. A two-way, mixed (SOA within, STREAM between) analysis of variance (ANOVA) yielded a significant first order interaction confirming a difference between the two stream types across SOA, $F(7,98) = 3.127$, $p < .05$. Similarly, significant main effects were observed for STREAM, $F(1,14) = 6.125$, $p < .05$, and SOA, $F(7,98) = 4.903$, $p < .05$.

Repeated measures ANOVAS were conducted on the AB magnitude data for each SOA for both the scrambled and smooth stream types to ascertain whether a

significant AB was present in each condition. As predicted by the pattern of data in

Figure 8, these analyses confirmed there was no effect of SOA in the smooth morph

stream where object file continuity is preserved, $F(7,49) = 0.500$, $p > .05$, but a

significant effect of SOA in the scrambled stream where object file continuity is

disrupted, $F(7,49) = 6.748$, $p < .05$. Furthermore, Figure 8 suggests that this AB

effect persisted until 267 ms. To investigate this, single sample t-tests (one tailed)

were carried out for the first four SOA's and the alpha level was set to .01, a

conservative level to reflect the number of comparisons made. These tests

demonstrated that the AB effect in the scrambled morph condition is significant at 67

ms, $t(7) = 4.244$, $p < .01$, at 133 ms, $t(7) = 4.274$, $p < .01$, and 200 ms, $t(7) = 4.782$,

$p < .01$, but has recovered by 267 ms, $t(7) = 1.220$, $p > .01$.

Figure 8.
Mean AB magnitude for each SOA for smooth and scrambled morph streams.
Standard error ($\pm$ 1) is represented by the vertical bars.

In the scrambled morph condition, despite the inherent disruption in object file continuity, each successive frame is generally similar in terms of low level spatial information. It is possible that the observed AB effect results from the updating of object file information rather than the generation of new object files - a process which is thought to necessitate greater attentional resources (Kahneman et al., 1992). One might therefore anticipate a larger AB effect (i.e. greater magnitude or longer temporal duration) where the stream requires the generation of new object files. One way to disrupt object file continuity and create the demand for new object files is to embed the same targets within a consecutive stream of random objects having intrinsically different spatial properties to the smoking pipe or the saucepan. This is examined in Experiment 2.

Experiment 2 – Random object RSVP stream

The experimental methods and procedure adopted were the same as those outlined in Experiment 1 with the following exceptions.

Design

A single factor, SOA (all eight levels) was employed in this experiment.

Participants

A total of eight participants were recruited. The sample was comprised of volunteers from the student participant panel (1 male, 3 females, mean age 23.7

years, SD = 4.65 years) and the community participant database (1 male, 3 females, mean age 23.5 years, SD = 6.40 years).

Stimuli and procedure

In order to create RSVP streams for this experiment the pipepan targets were embedded within a sequence of 22 pictures of random objects. An example stream is presented in Figure 9. As before, T1 appeared at position 12 in the stream and T2 could appear at one of the eight serial positions following T1.
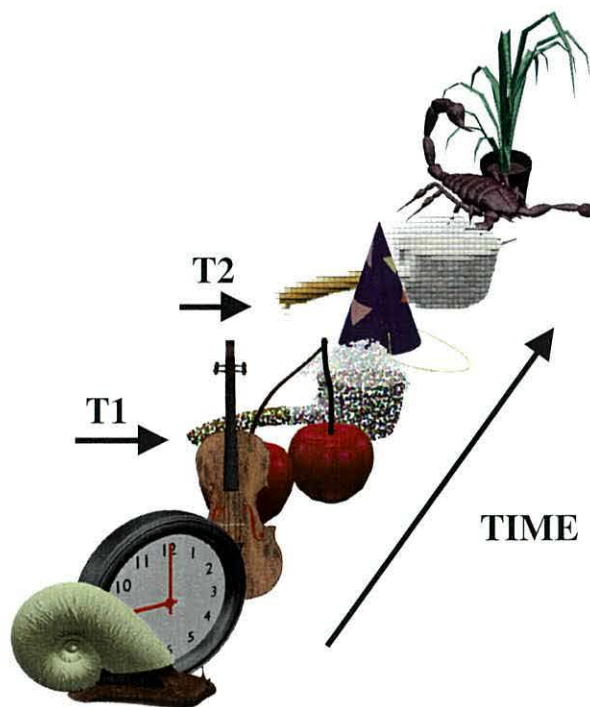
Figure 9.
Example random object RSVP stream. Streams were 24 frames in length but have been truncated here for practical purposes. T2 is shown in the second T2 serial position (corresponding to 133 ms).

Results and discussion

The mean percentage of correct T1 responses in the dual task trials was 87% (SD 12.07%). Mean AB magnitude was calculated for each SOA and these are shown in Figure 10 which suggests that an AB effect was obtained using a random object RSVP stream. This was confirmed by an ANOVA which revealed a significant main effect of SOA, $F(7,49) = 2.569$, $p < .05$.
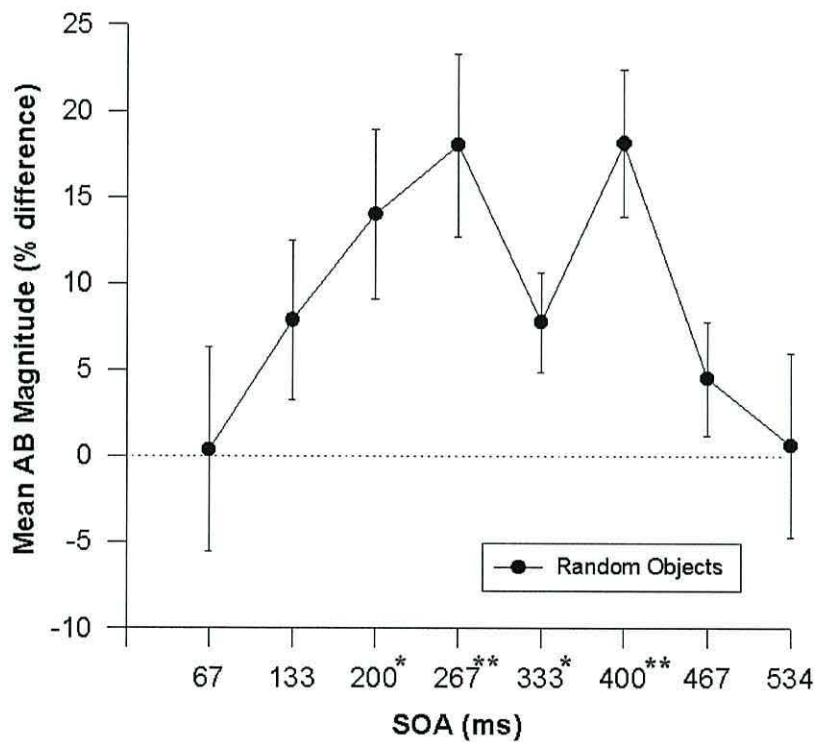
Figure 10.
Mean AB magnitude for random object stream. Standard error ($\pm$ 1) is represented by the vertical bars.

In order to establish the duration of the AB for the random object stream single sample t-tests were conducted for each of the AB magnitudes between 133 ms and 400 ms, using a .01 alpha level to reflect the number of comparisons being made. Because of the increased variance (compared to the scrambled morph stream in Experiment 1) and the modified alpha level, these analyses were only significant for 267 ms, $t(7) = 3.380$, $p < .01$, and 400 ms, $t(7) = 4.241$, $p < .01$. This demonstrates that the random object stream yielded an AB effect that did not recover until 467 ms. Using this alpha level the AB magnitudes at the SOA's were not significant; 133 ms, $t(7) = 1.699$, $p > .01$; 200 ms, $t(7) = 2.856$, $p > .01$; 333 ms, $t(7) = 2.700$, $p > .01$. However, the single asterisks in Figure 10 demonstrate that if one adopts a .05 alpha level, but is aware of the risk of committing an Type I error, then the 200 ms and 333 ms SOA's would also have been significant.

This experiment attempted to maximally disrupt object file continuity by utilising random object stimuli that should have initiated the attentionally demanding process of generating new object files. A robust AB was observed and this effect was longer in duration than that obtained in the scrambled morph condition from Experiment 1. It is interesting to note that the onset of the AB was not immediate - that is, lag-1 sparing was observed. This contrasts with the results for the scrambled morph in Experiment 1 where lag 1 sparing was not present.

General Discussion

The experiments in this chapter successfully created optimum conditions for the preservation or disruption of object file continuity. In the smooth morph condition where object file (and perceptual) continuity was preserved, complete attenuation of the AB was observed. In contrast, a robust AB effect was associated with object file (and perceptual) discontinuity that was created either by the scrambled morph or random object streams.

The scrambled morph stream yielded an AB for the first three serial positions. This corresponds to 267 ms and is a relatively short lived effect compared to other studies which report an AB that continues until 450 ms (e.g.. Raymond et al., 1992). This shortened effect may reflect object file updating rather than the generation of new object files. However, the AB duration here (267 ms) is longer than the 90 ms that Raymond and Sorensen (1994) associate with updating an object file.

It was anticipated that the object file discontinuity in the random object stream would demand the creation of new object files rather than the updating of existing ones. Since object file generation is thought to require greater attentional resources than amending an object file (Kahneman et al., 1992), one would predict a larger AB (magnitude or duration) for the random object stream compared to the scrambled morph. This is supported by the fact that the AB effect in the random object stream did not recover until 467 ms post T1 onset (200 ms after AB recovery in the scrambled morph condition).

Apart from differences in object file related factors, the RSVP streams used in Experiment 1 and 2 also differed in terms of the masks that followed T1 and T2. The importance of this factor in terms of the experimental outcomes is considered in detail in subsequent chapters.

Generally, one of two different AB functions is observed, a U–shaped function or a monotonic function. A review of the recent literature found that both patterns are reported with roughly equal frequency (Visser, Bischof & Di Lollo, 1999). The scrambled morph condition produced a monotonic function where the deficit is most pronounced at the first SOA, reducing as a function of increasing SOA. In contrast, the random object morph produced a U-shaped function (if we were to plot single and dual task performance instead of AB magnitude) where the onset of the AB is not immediate (i.e. there is lag-1 sparing) but is greatest for the second or third SOA.

Both Chun and Potter (1995) and Shapiro and Raymond (1994) postulate that lag-1 sparing occurs because of a slow closing attentional gate. In a similar vein, Sperling and Weichselgartner (1995) propose the concept of a discrete attentional episode or window of attention that can take up to 200 ms to close. By this account, lag-1 sparing occurs when both targets form part of a single attentional window (or episode). However, if this is the case then one might expect it to be obtained consistently (Visser et al., 1999).

One possibility is that lag-1 sparing is an indication of attentional switching; that is, how efficiently the system can, after processing T1, reconfigure itself in order to accommodate the different processing demands of T2 when it occurs in close

temporal proximity (Visser et al., 1999). When attentional switching is successful there is no deficit in T2 performance and therefore lag-1 sparing occurs. However, when attentional switching is unsuccessful (possibly because the reconfiguration process was unable to cope with excessive switching demands) T2 performance is impaired and lag-1 sparing is not obtained.

In a review of over 20 publications, Visser et al., (1999) attempt to identify those attentional switching conditions associated with the presence or absence of lag-1 sparing. They conclude that lag-1 sparing is obtained when there is no attentional switch between targets or when there is a unidimensional switch between targets, such as a switch in task (e.g. identify T1, detect T2) or category (e.g. T1 letter, T2 digit). In contrast, when there is a switch in location (T1 and T2 appear in different spatial locations) or a multidimensional switch between targets is required lag-1 sparing does not occur. This explanation does not explain the pattern of results for the scrambled morph and random object streams. Both streams involve a unidimensional switch in category (from dots to squares) and one might predict lag-1 sparing in both conditions, yet it was only present in the random object stream. In an attempt to explain this disparity one needs to consider the nature of the distractor items in each stream. In the random object stream, a perceptual filter for the pipepan stimulus will select T1. Since the item immediately following T1 (the +1 item) is usually a <u>non</u> pipepan item (which will not fit the filter) it is not essential for the filter to close rapidly in order to limit interference from the +1 item. Because the filter is slow to close T2 is opportunely processed (by virtue of sharing the same attentional window) and lag-1 sparing is obtained. In the scrambled morph, where all the stream items are similar, the same perceptual filter will need to close rapidly

in order for T1 to be processed efficiently whilst minimising interference from very similar +1 items. Because the filter closes rapidly T2 is unlikely to gain entry, and lag-1 sparing would not occur. Thus, perhaps the similarity of distractor items also plays a role in determining the pattern of the AB, and in particular the presence or absence of lag-1 sparing.

The results of these experiments generally agree with the work of Raymond and Sorensen (1994) and Kahneman et al. (1992). The first item in the RSVP stream creates an object file which is reviewed for each subsequent visual event. Where T1 is the same object as the preceding item (i.e. in the smooth morph stream) it can easily utilise the current object file. Because attentionally demanding object file operations have not been required there are sufficient resources to fully process T2 and consequently no AB is observed. In contrast, when T1 cannot adopt the current object file (i.e. in the scrambled morph or random object streams) attentionally demanding object file processes are required. These processes interfere with the processing of T2 at early serial positions and an AB is observed. Whilst the results of these experiments are in agreement with this conceptual model they do not (nor were they designed to) isolate the locus of the effect as being specific to the object file relationship between T1 and the preceding item as Raymond and Sorensen (1994) suggested. Alternative candidates for the locus of this effect include the object file relationship between the two targets and the nature of the items between the two targets. Chapter 3 investigates whether the process of morphing was crucial in providing object file continuity between the two targets in the smooth morph condition, or whether it is simply that both targets can share the same object file so

long as the information is not overwritten by intervening objects as in the random object stream. The role of the pre-T1 items is also investigated.

Before moving on to the next chapter it is important to briefly mention the nature of the streams used in the next two chapters. The next two chapters consider what factors modulate the AB where object continuity is preserved and where it is not. The part-stream investigations presented in Chapters 3 and 4 use the random object stream and the smooth pipepan morph. Their selection is based on the assumption that the morph stream represents the use of a single object file whereas the random object stream does not. Such a presupposition is based on the results presented in this chapter. The scrambled morph condition, which presumably involves updating an object file, does not feature in these investigations.

Chapter 3 – Investigating factors that modulate the AB in the smooth morph and random object streams.

Abstract

The three experiments in this chapter investigate factors which might have modulated the AB effect in the full stream morph (where no effect was found) and the random object condition (where a robust AB was observed). Experiment 3 considered whether morphing was crucial in providing object file continuity in the smooth morph condition and used the *critical items* paradigm (Duncan et al., 1994) whereby each stream was comprised of T1, T2 and their respective masks. An AB effect was not observed although this outcome is in agreement with the theoretical notion of 'recency' in object file reviewing. Experiment 4 tested this notion of recency and attempted to overwrite T1 object file information by presenting random objects between T1's mask and T2 such that they could not share the same object file. Once again the AB was attenuated, contrary to theoretical predictions. Experiment 5 sought to determine the role of the items preceding T1 and involved presenting the critical items and also the pre-T1 items from the random object and morph conditions. The work of Raymond and Sorensen (1994) would predict an AB effect in the random object condition but not in the morph condition. However, neither condition yielded an AB. The unexpected outcome for the part stream random object experiments may be due to the fact that a pipepan mask was used rather than a random object mask as used in Experiment 2 – this issue is addressed in Chapter 4 and Chapter 5.

Investigating factors that modulate the AB in the smooth morph

and random object streams.

The three experiments in this chapter attempted to identify factors that

modulated the AB effect in the smooth morph stream (Exp. 1) where the AB was

attenuated and the random object stream (Exp. 2) where a robust effect was obtained.

The decision to investigate these two stream types reflects the assumption that they

represent the use of the same object file (morph condition) or the generation of new

object files (random object condition). Experiment 3 examined the smooth morph

stream and investigated whether the processing of morphing was essential in

providing object file continuity. The other two experiments, which examined the

random object stream, considered the role of the items intervening T1 and T2 (Exp. 4)

and the items preceding T1 (Exp. 5).

Before moving on to these experiments it is instructive to consider the nature of

perceptual masking in the random object and smooth morph streams. To

recapitulate, the target in the random object stream was masked by a random object

whilst the same target in the smooth morph stream was masked by the next

consecutive item in the morph movie. Given that T1 must be adequately masked if

an AB is to be obtained (Raymond et al., 1992) it is possible that the outcomes

reflect inadequate masking in the smooth morph condition but efficient masking in the

random object condition. The argument is less convincing in relation to the smooth

and scrambled streams where different outcomes were obtained using almost identical

masks. To investigate this issue eight participants completed 480 trials in which they identified the pipepan T1 target (presented using the same parameters as in the full stream experiments). The singly presented target was either masked or unmasked. When the target was masked it was either by a random object or by the next consecutive item from the pipepan morph stream (morph frame 13). The results showed that when the mask was present participants reached a mean performance of 93% (SD 6.92%) in the morph condition and 94% (SD 9.27%) in the random object condition. When the mask was absent, performance in the morph condition remained at 93% (SD 10.02%) for the morph condition and 95% (SD 6.80%) in the random object condition. These results demonstrate that target performance was equivalent (and good) irrespective of the type of mask used or whether the mask was present. Since these results were suggestive of ceiling effects it was thought prudent to equate perceptual masking by using the smooth pipepan mask (next consecutive item) for all the part stream experiments presented in this chapter.

Experiment 3 – Taking the morph out of the morph stream

This experiment investigated whether the AB attenuation observed in Experiment 1 was modulated by the morphing process that provided object file continuity between T1 and T2. If so, one would expect the AB effect to return if the

morphing process were removed and the items intervening the two targets replaced

by a variable blank interval.

Design

This experiment used one factor: the seven SOA's of T2. Since T1's mask is

crucial to the critical items paradigm (see stimuli and procedure below) the first SOA

(67 ms) was excluded.

Participants

Twelve volunteers (5 female, 7 male, mean age 22.8 years, SD = 4.15 years) were

recruited from the community participant panel.

Stimuli & procedure

A *critical items* paradigm as used by Duncan et al., (1994) was employed. This

modified RSVP paradigm involves reducing the full pipepan stream to T1, T2 and their

respective masks and manipulating the SOA by inserting a blank interval between T1's

mask and T2. The mask item was the next consecutive item from the morph movie. A

typical RSVP stream is illustrated in Figure 11.

Due to a parameterisation error half of the participants completed 16 trials at

each SOA rather than 20 (112 trials in total).

TIME

VARIABLE TARGET-TO-TARGET ONSET

Figure 11.
The critical items RSVP stream involved presenting T1, T2 and their respective masks (which were the next consecutive item from the morph movie). Variable target-to-target onset was provided by inserting blank frames between T1's mask and T2. As before, participants were required to identify the size of the dots (for T1) and the size of the squares (for T2).

Results and Discussion

For the dual task trials the percentage of correct T1 responses was 83% (SD 12.23%). This was lower and more variable than T1 performance in the smooth morph condition and was comparable with performance in the full stream random

object condition. Mean AB magnitudes were computed for each SOA and these are displayed in Figure 12. A significant AB effect was not obtained, $F(6,66) = 0.750$, $p > .05$. Similarly these results were not found to be significantly different from the full stream smooth morph condition, $F(6,108) = 0.797$, $p > .05$.

·Figure 12.
Mean AB magnitudes for the critical items stream which involved the presentation of T1, T2 and their respective masks. Vertical bars represent standard error (+ 1).

These results suggest that in the smooth morph full stream (in Experiment 1) the morph items between the two targets were not crucial for providing object file continuity. It appears that a single object file can be used so long as the visual information is not inconsistent with the object file reviewing process but the greater

degree of variability in this experiment (compared to the morph condition in Experiment 1) suggests that the morphing process does contribute to perceptual continuity.

The outcome of this experiment can be explained in terms of object file theory which holds that when reviewing the current object file only the most recent visual information is used (Kahneman et al., 1992). When deciding whether the current visual display fits with the current object file, the visual information is compared with that immediately preceding it. Therefore, in this experiment it is difficult to establish whether the T1/T2 intervening pipepan items were necessary to provide T1/T2 continuity because the frames which were displayed were consistent with the current object file information. Additionally, because the pipepan movie was 24 frames long there little object change over three or four frames - removing the intervening items does not constitute much of a mismatch between the displayed frames. With this in mind, Experiment 4 develops the notion of recency and attempts to overwrite the T1 object information with that of a different object so that T2 cannot use the object file without making substantial modifications.

Experiment 4 – Critical items plus random intervening objects – a test of recency in object file reviewing.

This experiment was the same as Experiment 3 with the following exceptions.

Participants

Eighteen volunteers (10 female, 8 male, mean age 23.5 years, SD = 5.48 years) were recruited from either the Community Participant Database (15) or the Student Participant Panel (3).

Stimuli & procedure

As in the previous experiment the critical items paradigm was used (i.e. T1, T2 and their respective masks) but rather than manipulate SOA by inserting blank frames, a number of random object pictures were inserted between T1's mask and T2. The number of random objects in each stream was dependent on SOA and ranged between zero and six. A typical RSVP stream is illustrated in Figure 13.

Due to a parameterisation error fourteen participants completed 16 trials at each SOA rather than 20 (112 trials in total).

TIME

VARIABLE TARGET-TO-
TARGET ONSET

Figure 13.
An example of the modified critical items streams used in Experiment 4.
Depending on the serial position of T2, zero to six random objects were
presented between T1's mask and T2. Figure shows T2 occurring in serial
position 4 which corresponded to an SOA of 267 ms.

Results and Discussion

The data from one participant was replaced because their performance

approached the exclusion criterion for all measures (T2 performance in the single task

trials was 55% and T1 performance in the dual task trials was 51% correct). A high

degree of participant variability was apparent in the data set. Therefore, the single

and dual task data were subjected to a boxplot analysis to identify outlying data

points. Where an outlier was identified, all that individual's data were removed. This

reduced the data set to seven participants. Subsequent analyses relate to this reduced

set.

The mean percentage of correct T1 responses in the dual task trials was 84%

(SD, 16.42%). Mean AB magnitudes were calculated and these are displayed in

Figure 14.

Figure 14.
Mean AB magnitudes for the modified critical items stream that involved
presenting a variable number of random objects between T1's mask and T2.
Vertical bars represent standard error (± 1).

Whilst Figure 14 is suggestive of an AB effect, an ANOVA conducted on the AB magnitudes showed there to be no significant effect of SOA, $F(6,36) = 1.02$, $p > .05$. However, the reduced sample size is likely to be a major factor in this outcome since sample size reduction results in reduced statistical power - the probability that a false null hypothesis can be correct rejected (Howell, 1997). Despite the absence of a significant overall AB there does appear to be a hint of an effect at the 200ms SOA which was subsequently shown to be significant, $t(6) = 2.853$, $p < .05$ (one tailed). It is instructive at this point to consider the nature of the stream at the 133 ms SOA. Streams representing the 133 ms SOA were comprised of four consecutive pipepan items with no intervening random objects. Therefore, it is not surprising that a non-significant effect was obtained for this SOA because there were no random objects to overwrite the object file information and object file continuity was preserved.

Theoretically, a significant AB was anticipated. Drawing on the notion of recency, the object file information generated for T1 should be overwritten by information relating to the objects intervening T1 and T2. Therefore, when T2 was presented, the process of reviewing the current object file would involve a mismatch in visual information (a pipepan compared to a random object) and the object file would need to be modified in order to accommodate T2. Whilst the significant outcome at the 200 ms ISI is consistent with such a notion, an AB effect was not observed overall. Additionally, these data were not significantly different from the previous experiment which used a blank interval to vary SOA rather than random objects, $F(6,102) = 1.432$, $p = 0.210$. However, the notion of recency in object file

reviewing is supported by the outcomes for the first two SOAs and based on this outcome one can conclude that the stream items intervening T1 and T2 may be implicated in object file continuity and AB modulation. An alternative theoretical explanation for the outcomes in the random object and smooth morph full streams is based on the items preceding T1 as in Raymond and Sorensen (1994). This is examined in Experiment 5.

Experiment 5 – Critical items plus pre-T1 items – investigating the role of the items preceding T1.

This experiment investigated the role of items preceding T1 in the full stream morph and random object conditions. The rationale was based on notion that the first stream item creates an object file that subsequent items, including T1, may or may not be able to adopt (Raymond and Sorensen, 1994). When subsequent items are not the same object (as in the random object stream) the object file must be modified, or a new one generated, in order to accommodate T1. Because this process is attentionally demanding, T2 processing should experience interference from the object file operations being undertaken for T1. Therefore, one might anticipate an AB effect in the random object condition but not in the morph condition. This experiment was the same as Experiment 3 with the following exceptions.

Design

This experiment employed a 2 x 7 mixed factorial design. The between groups factor (STREAM) related to the nature of the items preceding T1 and comprised two levels: MORPH where the pre-T1 items were the first 11 items from the morph movie and RANDOM where the pre-T1 items were 11 random object pictures. The within groups factor was the seven T2 SOAs.

Participants

Twenty volunteers were recruited from either the Community Participant Database (6) or the Student Participant Panel (14). Participants were allocated randomly to either the morph (3 male, 7 female, mean age 22.4 years, SD = 3.38 years) or random object condition (9 female, 1 male, mean age 22.6 years, SD = 5.61 years).

Stimuli & procedure

This experiment used the same stimuli as Experiment 3 where T1, T2 and their respective masks were presented with a variable blank interval creating the SOA between T1's mask and T2. However, in this experiment the critical items paradigm was modified by the addition of eleven pre-T1 items from the full stream morph or random object picture streams. Example streams are shown in Figure 15.

TIME

VARIABLE TARGET-
TO TARGET ONSET

MORPH   RANDOM OBJECTS

Figure 15.
An example of the modified critical items streams used in Experiment 5.
Eleven items (random objects or sequential items from the morph
movie) were presented prior to T1. As in Experiment 3 the SOA was
provided by inserting a variable number of blank frames between T1's
mask and T2.

Results and Discussion

Data from three participants in the random object group were replaced (one

particpant had a 66% error rate in the single task trials and the others had 58% and

51% T1 errors in the dual task trials). T1 performance was at 96% (SD = 13.87%)

for the morph group and this was significantly higher than T1 performance for the random object group which was at 78% (SD = 16.64%), $t(138) = 6.677$, $p < .01$. Mean AB magnitudes were calculated for the morph and control groups. These are shown in Figure 16 which suggests that, at least for the random object group, there may be a SOA effect for the first three SOAs. However, a mixed (STREAM between, SOA within), two-way ANOVA revealed a non-significant first order interaction $F(6,108) = 0.877$, $p > .05$ and non-significant main effects of stream, $F(1,18) = 0.263$, $p > .05$ and SOA, $F(6,108) = 1.299$, $p > .05$. One must conclude that there was no significant difference in AB magnitudes across all SOAs and between both stream types. Additionally, the data from Experiment 3 (critical items alone) were not found to be significantly different from either the morph, $F(6,120) = 0.498$, $p > .05$, or random object conditions, $F(6,120) = 1.147$, $p > .05$ in this experiment.

For exploratory purposes the following individual ANOVAs were conducted. These yielded non-significant effects of SOA for both the smooth stream, $F(6,54) = 0.822$, $p > .05$ and the random object stream, $F(6,54) = 1.238$, $p > .05$. The pattern of data shown in Figure 16 prompted t-test analyses (one tailed) at specific SOAs for each stream type. As one might expect, this yielded a significant effect for the random object group at the 133 ms SOA, $t(9) = 2.596$, $p < .05$ and the 400 ms SOA, $t(9) = 1.855$, $p < .05$ but not for the 200 ms SOA, $t(9) = 1.588$, $p > .05$. A single analysis was conducted for the morph group (for the 133 ms SOA) and this was significant, $t(9) = 2.008$, $p < .05$.
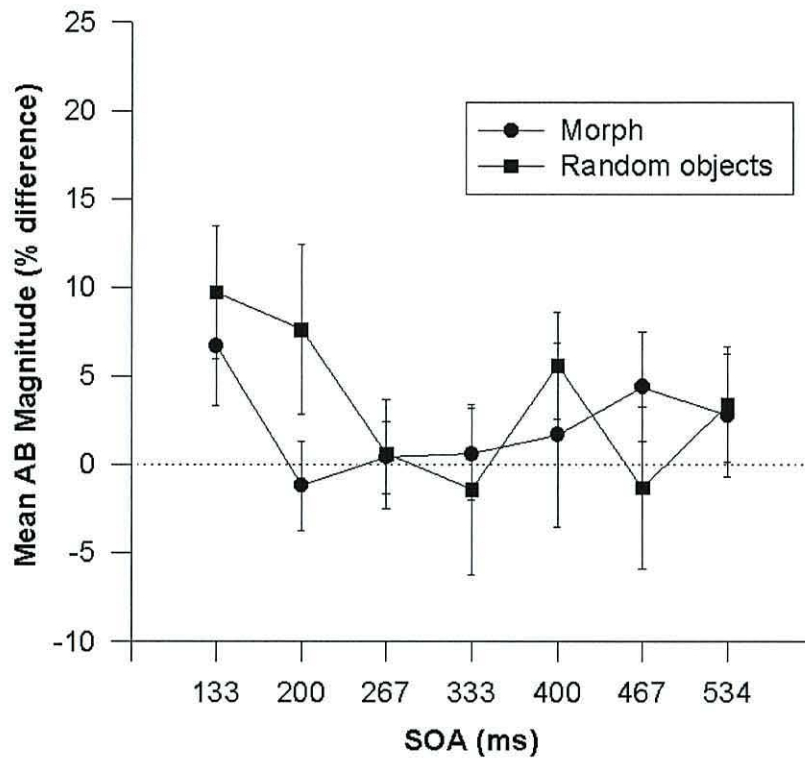
Figure 16.
Mean AB magnitudes for the modified critical items stream that involved presenting the pre-T1 items from either the smooth morph stream or random object stream. Vertical bars represent standard error (± 1).

It was anticipated that a significant AB would be obtained for the random object stream but not for the morph group. However, a significant AB effect was not observed in either condition. In terms of Raymond and Sorensen's (1994) theoretical framework, the morph stream should not yield an AB because each stream item can use the object file created by the initial item. The random object stream, on the other hand, should have yielded an AB because the pipepan T1 should not have been able to use the object file created by a different object.

General Discussion

The three experiments in this chapter investigated which aspects of the full stream morph and random object conditions (from Chapter 2) modulated the AB effect. Experiment 3 examined whether, in the full stream morph condition, the process of morphing was crucial in providing object file continuity. Despite removing the smooth process of object change associated with the full stream no AB was observed, thus demonstrating that the process of morphing was not crucial in providing object file continuity; increased variability here suggests that the morphing process did assist in providing object file continuity. The results are in theoretical agreement with the object file notion of recency which maintains that only the most recent visual information is used during the object file review process (Kahneman et al., 1992).

The notion of recency featured in Experiment 4 which involved the same critical items but with random objects presented between T1's mask and T2. Theoretically, if T1 object file information is overwritten by a different object then T2 should not be able to easily adopt the object file. Despite anticipating a significant AB, it was not obtained and these results were not significantly different to those obtained for critical items alone (Experiment 3). The absence of an overall AB effect in this experiment does not accord with Kahneman et al.'s (1992) notion of recency in object file reviewing which would predict a robust effect. However, the non-significant effect for the 133 ms SOA (where no random objects were presented) and the significant effect at the 200 ms SOA do support such a notion, and suggests that

the intervening items in the random object full stream may have been a contributory factor in object file continuity and AB modulation.

Experiment 5 investigated whether the pre-T1 items are important in AB modulation. Here, the critical items were presented with the pre-T1 items from either the morph or random object stream. Experimental predictions were based on the work of Raymond and Sorensen (1994). In the smooth morph condition, where T1 should be able to use the object file set up by the first stream item, an AB was not predicted and was not obtained (apart from a significant effect at the 133 ms SOA). An AB effect was predicted in the random object stream because T1 is a different object to all its preceding items, although a significant effect was not obtained overall. Additionally, both the morph and random object data were not significantly different from the critical items alone (Experiment 3). However, given that the 133 ms and 400 ms SOA's yielded significant effects for the random object stream one could claim that the pre-T1 items may also have a role to play in AB modulation. Surprisingly, there was also a small but significant effect in the morph condition at the 133 ms SOA. This could be explained in terms of the participants employing a strategy based on the gap. For instance, for all but the 133 ms SOA the appearance of T2 is cued by a blank interval. However, if one uses this strategy as the basis for attentional monitoring then targets occurring at the 133 ms are less likely to be identified correctly. Such an explanation could also be applied to the significant 133 ms SOA effect in the random object condition but not for the 400 ms SOA.

No complete answer emerges from these experiments as to what aspect of the morph and random object full streams modulated the effects that were obtained (in

Chapter 2). The differences observed at specific SOA's in Experiments 4 and 5 are supportive of the notion that the effects are additive (i.e. small effects resulting from intervening items and pre-T1 items combine to yield an overall effect) and that it takes the whole random object stream in order to obtain the effects observed in Experiment 2. Before asserting such a claim it is important to consider particular features of the random object experiments in this chapter which may have influenced the experimental outcomes. One important feature relates to masking. To recapitulate, perceptual masking was investigated in order to ascertain whether a pipepan target is masked equally well by a random object or another pipepan item. Whilst the results were not significantly different it was thought prudent at the outset to equate masking in this chapter by using the same mask across all condtions. Therefore, the part stream experiments in this chapter used pipepan masks whereas the full stream random object experiment (Experiment 2) used a random object mask. The outcomes for Experiments 4 and 5 may be due to the use of a pipepan mask rather than a random object mask. For this reason Chapter 4 re-examines the part stream random object experiments and investigates the role of the items intervening the two targets and the pre-T1 items by conducting part stream experiments using a random object mask.

Chapter 4 – random object part stream experiments using

a random object mask

Abstract

The experiments in this chapter are replications of two experiments from the previous chapter. Experiment 6 examined the role of the random objects intervening T1 and T2 but used a random object mask (to match masking in Experiment 2, the full stream of random objects). A significant AB was obtained (at 133 ms and 267 ms) suggesting that the items intervening T1 and T2 are involved in modulation of the AB effect in the full stream of random objects. These results agree with Kahneman et al.'s (1992) notion of recency in object file reviewing: T1's object file information would be overwritten by the random objects appearing between T1 and T2 such that they could not share the same object file. In Experiment 7 the role of the pre-T1 items was examined but using a random object mask. Despite anticipating a significant AB effect, it was not obtained. Such an outcome disagrees with the work of Raymond and Sorensen (1994), neither does it agree with Kahneman et al.'s (1992) notion of recency since the random object mask should overwrite the object file information for T1. If the locus of the AB modulation in the full stream experiment was due to the random object mask a significant AB effect would be predicted in both experiments. Given that only Experiment 6 yielded such an effect one can conclude that AB modulation in the full stream was due to a combination of the random object mask and the random objects intervening T1 and T2.

Random object part stream experiments using

a random object mask.

The two experiments in this chapter are replications of two experiments from Chapter 3. However, these experiments used a random object mask rather than a pipepan mask. Experiment 6 was a replication of Experiment 4 (critical items with random objects between T1 and T2) and Experiment 7 was a replication of the random object condition in Experiment 5 (critical items with pre-T1 items).

Experiment 6 – Critical items plus random intervening objects and a

random object mask.

This experiment was the same as Experiment 4 with the following exceptions.

Participants

Sixteen volunteers (12 female, 4 male, mean age 21.6 years, SD = 3.76 years) were recruited from either the Community Participant Database (3) or the Student Participant Panel (13).

Stimuli & procedure

RSVP streams were the same as those used for Experiment 4 with the exception that random objects were used to mask T1 and T2 rather than pipepan stimuli. The

number of random objects appearing in each stream ranged between one and seven (depending on SOA). A typical RSVP stream is illustrated in Figure 17.

Figure 17.
An example of the modified critical items streams used in Experiment 6. Apart from the random object mask, streams were the same as those used in Experiment 4. Depending on the serial position of T2, one to seven random objects were presented between T1 and T2. Figure shows T2 occurring in serial position 4 which corresponds to a 267 ms.

Half of the participants completed their two sets of experimental trials in two sets: set 1 consisted of 56 trials (yielding 8 trials at each SOA) and set 2 comprised 84 trials (yielding 12 trials at each SOA). By splitting the trials into sets it was

possible to establish whether any overall effects were affected by increased

variability in the first part of the experimental trials (possibly due to unfamiliarity

with the task). The remaining 8 participants completed two sets of 140 trials (20

trials at each SOA) as in previous experiments. In this case, because of the way in

which the morph program works, it was not possible to analyse the single set data in

terms of first or second halves. This is because the program randomly selects from

the total number of trials (140) rather than the total stream set (28) and consequently

an unequal number of trials at each SOA would be obtained if one were to split the

data into two sets based on order of presentation.

Results and Discussion

Data from two participants were replaced because they made excessive T2 errors

(54% and 58%) in the single task trials.

Before pooling the data, data from participants who completed two sets of trials

was analysed. Set 1 data appeared to be more variable than set 2 data and an

ANOVA confirmed a significant main effect of set, $F(1,7) = 7.981$, $p < .05$. This

suggests that during the first set, despite having completed 56 practice trials,

participants were still becoming accustomed to the task. Consequently, set 1 data

were removed and only set 2 data were pooled with the remaining data (where trials

were completed in a single set). This reduced data set was used for all subsequent

analyses.

The mean percentage of correct T1 responses in the dual task trials was 84%

(SD, 17.32%). Mean AB magnitudes were calculated and these are displayed in

Figure 18.

Figure 18.
Mean AB magnitudes for the modified critical items stream that involved
presenting a variable number of random objects between T1's mask and T2 as
well as using a random object mask. Vertical bars represent standard error ($\pm$ 1).

Figure 18 is suggestive of an AB effect. This was confirmed by an ANOVA

which revealed a significant effect of SOA, $F(6,90) = 2.25$, $p < .05$. These results are

not significantly different to the original full stream of random objects, $F(6,132) =$

1.175, p > .05 which also yielded a significant AB. Single sample t-tests (one-tailed) were conducted at all SOA's except 400 ms and 534 ms (using a .01 alpha level). These analyses were significant for 133 ms, $t(15) = 3.182$, $p < .01$, and 267 ms, $t(15) = 4.654$, $p < .01$. However, non-significant results were obtained for AB magnitudes at SOA's of 200 ms, $t(15) = 1.469$, $p > .01$; 333 ms, $t(15) = 1.489$, $p > .01$; and 467 ms, $t(15) = 1.634$, $p > .01$. If one were using an alpha level of .05 these three SOA's would also approach significance.

These results are consistent with Raymond and Sorensen's (1994) theoretical framework. Since no items preceded T1 one would expect the presentation of T1 to initiate an open file process for each trial. This process consumes processing resources to the extent that there are insufficient resources to enable T2 to be processed to the level of report and an AB effect would be anticipated.

In terms of Kahneman et al.'s (1992) object file theory, if the pipepan information for T1 is overwritten by the random objects intervening T1 and T2 then the object file will need to be modified in order to accommodate T2 and one might predict an AB. The results are consistent with this notion of recency in object file reviewing and suggests that the locus of AB modulation in the full stream of random objects (Experiment 2) was either the items intervening the two targets, the random object masks, or a combination of the two. A comparison between this experiment (which used a random object mask) and Experiment 4 (which used a pipepan mask) shows that the results are not significantly different, $F(6,126) = 1.320$, $p > .05$.

Therefore, an explanation based solely on the random object mask becomes less convincing. In contrast, the significant overall AB for this experiment coupled with the significant effect (at the 200ms SOA) in Experiment 4 lends support to an explanation involving the intervening items.

Experiment 7 examines the role of the items preceding T1. It also uses a random object mask. If the random object mask is sufficient to have caused an AB in the full stream condition one might anticipate a significant effect there as well.

Experiment 7 – Critical items plus pre-T1 random objects and a random object mask.

This experiment was the same as Experiment 5 with the following exceptions.

Participants

Sixteen volunteers (11 female, 5 male, mean age 21.5 years, SD = 2.63 years) were recruited from either the Community Participant Database (11) or the Student Participant Panel (5).

Stimuli & procedure

RSVP streams were the same as those used for experiment 5 but used a random object to mask T1 and T2. A typical RSVP stream is illustrated in Figure 19.

As in the previous experiment half of the participants completed their single and dual experimental trials in two sets and the remaining 8 participants completed two sets of 140 trials (20 trials at each SOA).

TIME

VARIABLE TARGET-
TO-TARGET ONSET

Figure 19.
An example of the modified critical items streams used in Experiment 7. Apart from the random object mask, streams were the same as those used in Experiment 5 (random object condition). Eleven random objects always preceded T1. A variable blank interval (depending on SOA) occurred between T1's mask and T2.

Results and Discussion

Data from one participant was replaced because they made 51% T2 errors in the single task condition.

As in the previous experiment, prior to pooling all data, data from participants who completed two sets of trials was analysed. This analysis revealed that the data for set 1 and set 2 were not significantly different, $F(1,7) = 3.343$, $p > .05$. However, for consistency with Experiment 6 only set 2 data was pooled with the single set data and this reduced data set was used in all subsequent analyses.

The mean percentage of correct T1 responses in the dual task trials was 85% (SD, 14.58%). Mean AB magnitudes were calculated and these are displayed in Figure 20 which suggests there to be no AB. This was confirmed by a non significant effect of SOA, $F(6,90) = 1.591$, $p > .05$. Interestingly, five participants exhibited large order effects in that their single task performance was worse than their dual task performance (yielding a number of negative values for nearly all SOA's). Generally, counterbalancing the order of the single and dual task trials will control for adverse effects of task order. Therefore, such an outcome may have been a function of the stream type. Given that all individuals begin with three sets of dual task practice trials it may have been difficult for these participants to subsequently ignore T1 if they have just completed either the dual task trials (one participant) or the three sets of dual task practice trials (four participants). To successfully achieve this outcome

participants would need to attend to the first visual event that occurs after a variable

blank interval (because there are no items presented between T1's mask and T2) and

it may be harder to attend to (or ignore) a variable blank offset than a number of

random objects.

Figure 20.
Mean AB magnitudes for the modified critical items stream that involved
presenting eleven random objects prior to T1 as well as using a random object
mask. Vertical bars represent standard error ($\pm$ 1).

The results of this experiment also disagree with the work of Raymond and

Sorensen (1994). Given that there were pre-T1 items present and T1 was always a

different object to the preceding items one would predict an AB (although Raymond

and Sorensen might anticipate a shortened AB because both T1 and T2 are the same

object). However, a significant AB was not observed, suggesting in the full stream of random objects (Experiment 2) the items preceding T1 played little, if any, role in modulating the AB effect.

A comparison between this experiment (which used a random object mask) and Experiment 5 (which used a pipepan mask) shows that the results are significantly different, $F(6,144) = 2.426$, $p < .05$, but the data from each experiment is inconsistent with the notion that a random object serves as a more effective mask than a pipepan. In the present experiment the random object mask was associated with a non-significant AB overall and no significant effects at individual SOA's within the time-course of the AB. In contrast, a pipepan mask (Experiment 5) did yield significant outcomes for the first two SOA's (133 ms and 200 ms) despite the absence of an overall AB.

General Discussion

The two experiments in this chapter replicated Experiments 4 and 5. As in the previous chapter, these experiments examined whether the locus of AB modulation in Experiment 2 was centred on the items that intervene T1 and T2 (Experiment 6) or the items that precede T1 (Experiment 7). Both experiments used a random object mask to match the original full stream random object experiment.

In Experiment 6 a significant AB was observed, although this effect was earlier than that obtained in the Experiment 2. Here, the 133 ms and 267 ms SOA's were significant (using a .01 alpha level), whereas in Experiment 2 significance was obtained at 267 ms and 400 ms. The results of this experiment suggest that in Experiment 2 the locus of AB modulation was due either to the random objects intervening T1 and T2, the random object mask, or a combination of the two. Nevertheless this outcome is in agreement with Kahmenan et al.'s (1992) notion of recency in object file reviewing.

In Experiment 7, the role of the items preceding T1 was examined. In contrast to the theoretical predictions based on the work of Raymond and Sorensen (1994) no AB effect was observed, suggesting that in Experiment 2 the locus of AB modulation was not due to the items preceding T1. The Kahneman et al. (1992) recency notion should apply here also because the T1 pipepan information should be overwritten by its random object mask. Although, it might be true the case that a single random object is insufficient to completely overwrite the T1 information, especially since it will have been fully processed on virtually every trial.

An alternative explanation of the results in Chapter 2 could be due to differential effects of perceptual masking. If a random objects are more efficient than pipepans at masking T1 then one might anticipate a robust AB in the full stream random object stream (Experiment 2) and an attenuated AB in the full stream morph

stream (Experiment 1). Therefore, in this chapter, one should have obtained significant effects in both experiments since a random object mask was present in each case. Given that only Experiment 6 yielded an AB effect one can conclude that the random objects intervening T1 and T2 served to modulate the AB effect in the full stream experiment. However, given that no overall AB was obtained from the same condition but using a pipepan mask (Experiment 4), one must conclude that the full stream effect was modulated by a combination of the random object mask and the items intervening T1 and T2. It seems that neither the random object mask, nor the items intervening T1 and T2 are sufficient individually to modulate the AB effect.

Chapter 5 is the last of the experimental chapters. In that chapter the same pipepan mask is utilised in three full stream conditions: smooth morph, scrambled morph and random objects.

Chapter 5 – equating targets and masks across three full stream conditions

Abstract

This final experimental chapter is comprised of a single full stream experiment (Experiment 8) in which the nature of the target masks was controlled for all three RSVP stream types: smooth morph, scrambled morph and random objects. A texture pipepan mask – comprising of both big and small squares and dots - was used instead of a non-textured object (as employed in Experiments 1 and 2). The smooth morph condition was characterised by AB attenuation and this cannot be attributed to inadequate masking because T1 performance was decreased compared to Experiment 1 where a non-texture mask was used. The scrambled morph stream resulted in a marginally significant AB effect but a significant AB effect was present in the random object stream and this did not recover until 640 ms. Stream by stream comparisons demonstrate that these results are not significantly different to those obtained for the original full stream experiments reported in Chapter 2. This experiment confirms that the results of those original experiments were not due to differences in visual perceptual masking but to object file related factors. Furthermore, the fact that the scrambled and random object streams yielded different outcomes, despite having identical targets and masks, suggests that the shape and magnitude of the AB is not solely determined by those 'critical' items but is affected by other members of the RSVP stream and their object file relations.

Equating targets and masks across three full stream conditions.

The single experiment (Experiment 8) in this chapter marks a departure from critical items experiments and a return to the full stream RSVP paradigm. As in Chapter 2, this experiment employs streams comprised of smooth morph, scrambled morph or random object images.

In Chapter 2 it was claimed that the results of the original full stream experiments (Experiments 1 and 2) were due to differences in object file continuity. An alternative explanation relates to the efficacy of target masking. It is possible that a random object serves as a more efficient mask than the mask used in the smooth morph condition which was the next consecutive pipepan. If this were true then the outcomes – no AB in the morph condition but a significant AB in the random object condition – would be predicted. However, the same argument does not hold for the scrambled morph condition where a robust AB effect was obtained despite using a pipepan mask chosen randomly from the remaining morph images. With this in mind, the rationale underlying Experiment 8 was twofold. The first was to equate visual perceptual masking across all three streams and the second was to ensure that the textured targets were adequately masked by using a textured mask.

If the results of Experiment 1 and 2 were due to differences in perceptual masking then, when masking is equated across all three stream types, one might expect a robust AB effect of equal magnitudes in all conditions. Alternatively, if the

results of the original full stream experiments were due to object file related factors then one should observe similar outcomes to those obtained in Chapter 2. That is, a substantial AB in the random object stream which represents the generation of new object files, a reduced AB in the scrambled morph stream which represents the updating of object files, and an attenuated AB in the smooth morph stream in which a single object file should be utilised.

Experiment 8 – Full stream morph, scrambled morph and random objects with composite texture mask

This experiment was a replication of Experiments 1 and 2 (from Chapter 2) with the following exceptions.

Participants

36 volunteers (26 female, 10 male, mean age 23.3 years, SD = 4.83 years) were recruited from either the Community Participant Database (25) or the Student Participant Panel (11).

Stimuli & procedure

RSVP streams for this full stream experiments were the same as those presented in Chapter 2 (Experiments 1 and 2) but in this experiment the target's masks were

equated as much as possible across all types of stream: smooth morph, scrambled morph and random objects. All targets were masked by a composite texture mask. These were created by editing the underlying pipepan image such that it was comprised of a mixture of big and small dots as well as big and small squares. Morph movie frames numbered 13 to 21 were adapted in this way. Targets in the smooth morph condition were masked by a composite version of the next sequential pipepan image (i.e. where T2 was derived from the fifteenth pipepan image the mask for that target was derived from the sixteenth pipepan image).

The same composite texture masks (morph frames 13 to 21) were used to mask targets in the scrambled and random object condition with the proviso that the mask was never the next consecutive item (in order to avoid creating a short morph episode in either condition). Target-mask combinations were identical in the scrambled and random object condition (for example, in the scrambled morph condition, if T1 [frame 13] was masked by a composite version of morph frame 17 and T2 [frame 15] was masked by a composite frame 19 then the same target-mask combination was used for the equivalent stream in the random object condition). An example composite mask is shown in figure 21 and a typical RSVP stream is illustrated in Figure 22.

Figure 21.
An example of a composite pipepan mask which comprised both big and small squares and dots. Masks were created for eight morph frames (13 to 21).

Figure 22.
Example RSVP Streams for the smooth morph (left), scrambled morph (centre) and random object (right) conditions in Experiment 8. Targets in the smooth morph condition were masked by a textured version of the next consecutive frame. Targets were never masked by a consecutive image in the scrambled or random object conditions where masks were randomly selected from the eight texture masks. The choice of masks was identical in the scrambled morph and random object condition. Streams have been truncated for practical purposes.

As in previous experiments participants completed three sets of practice trials but at a slower rate. The first set comprised 16 trials (4 trials at each T2 serial position 5-8 corresponding to SOA's of 533 ms, 639 ms, 746 ms and 853 ms) at the slower rate of 186.62 ms per stimulus (5 items per second). The second set comprised 12 trials (4 trials at each T2 serial position 2-4 corresponding to SOA's of 213 ms, 320 ms and 427 ms) at the shorter duration of 159.96 ms per item (6 items per second). The third set of practice trials was comprised of 28 trials (4 trials at all SOA's) at a shorter duration of 106.64 ms per stimulus which corresponds to a presentation rate of 9 items per second.

140 experimental trials (20 trials for each of the seven serial positions of T2) were presented at a rate of 9 items per second. All participants completed their two sets of experimental trials (single and dual task trials) in two sets: set 1 consisted of 56 trials (yielding 8 trials at each SOA) and set 2 comprised 84 trials (yielding 12 trials at each SOA).

Results and Discussion

In the scrambled morph condition, data from three participants was replaced because they made excessive T2 errors (54%, and 66%) in the single task trials or

excessive T1 errors (54%) in the dual task trials. Similarly, in the random object condition, data from five participants was replaced because they made excessive T2 errors (50%, 50%, 54%, 58% and 61%) in the single task trials. In contrast, no data from the smooth morph condition needed to be replaced on the basis of T2 or T1 accuracy criteria.

Despite there being no significant difference of set across all stream types, $F(1,33) = 3.343$, $p > .05$, only data from the second set (84 trials, 12 trials at each ISI) was used in all subsequent analyses to provide consistency with previous experiments reported here (Experiment 6 and 7).

The mean percentage of correct T1 responses in the dual task trials was virtually identical for the smooth morph (mean 79%, SD, 14.60%), scrambled morph (mean 80%, SD, 16.09%) and random object (mean 78%, SD, 18.22%) stream types. On the basis of boxplot analyses, data for three participants was removed (one from each group) reducing the overall sample size from 36 to 33 (11 in each group). Mean AB magnitudes were calculated for each stream type and these are displayed in Figure 23.

Figure 23.
Mean AB magnitudes for the smooth morph, scrambled morph and random objects conditions. In this full stream experiment T1 and T2 were masked by composite texture masks. Vertical bars represent standard error (± 1).

A qualitative analysis of Figure 23 suggests there to be an effect of SOA in each condition, although to different degrees. As anticipated there seems to be a greater effect in the random object condition compared to the scrambled condition and a greater effect in the scrambled condition compared to the smooth morph condition. Given that each stream type yielded at least some effect of SOA it is not surprising that a non significant interaction was obtained between stream type and SOA,

$F(12,180) = 1.123$, p > .05.   The main effect of SOA was significant, $F(6,180) =$ 3.709, p < .00, but the main effect of stream was not, $F(2,30) = 2.005$, p > .05[3].

This significant effect of SOA was again evident in a trend analysis which confirmed there to be a global linear decrease in AB magnitude with increasing SOA, $F(1,30) = 19.209$, p < 00. However, in contrast to the overall ANOVA, the trend analysis revealed that the local changes in the slope of the linear trend were significantly different across the three stream types, $F(2,30) = 4.923$, p< .02. Trend analysis has been used previously in AB research, see for example, Brehaut, Enns and Di Lollo (1999); Shih, (2000); and Kawahara, Di Lollo and Enns (2001). In order to unpack the global linear trend, additional trend analyses were also conducted for each stream type. These will now be presented alongside stream by stream comparisons between this experiment and the original full stream experiments reported in Chapter 2. The role of the textured mask is then considered.

Random object stream

For the random object stream a significant linear trend was observed, $F(1,10) =$ 23.405, p < .02. Subsequent analyses (one-tailed t-tests) revealed significant AB magnitudes at 213 ms, $t(10) = 4.084$, p < .01; 320 ms, $t(10) = 4.921$, p < .01; and 533 ms, $t(10) = 3.418$, p < .01. A non-significant outcome was obtained for the 427 ms SOA, $t(10) = 1.211$, p > .05, this outcome was, however, due to an extreme

---

[3] A number of t-tests were conducted for Experiment 8 in order to make pairwise comparisons between streams at specific SOA's. However, these all non-significant (apart from smooth v RO at 213 ms) and are not reported here but are listed in Appendix B.

negative score and as soon as one individual with a score of –25 was removed the p

value for this analysis reached < .05 level. Therefore, for the random object stream a

robust AB effect was observed and this did not recover until 640 ms.

We have consistently seen the random object <u>full stream</u> result in an AB effect,

irrespective of whether targets were masked by a non-textured random object (as in

Experiment 2) or by a textured pipepan (Experiment 8). Furthermore an ANOVA[4]

shows there to be no statistical difference in AB magnitude across SOA between the

two experiments, $F(3,51) = 1.998$, $p > .10$. In Experiment 2 lag-1 sparing was

observed and discussed. In contrast, the significant linear trend in this experiment

suggests that lag-1 sparing was not present. Whilst the differences in timing

parameters make it impossible to make direct comparisons between the two

experiments, it is interesting to note that at around 200 ms (i.e. the third serial

position in Experiment 2 and the first serial position in Experiment 8) yielded

comparable AB magnitudes.

The scrambled morph stream

The trend analysis conducted on the scrambled morph stream approached

significance, $F(1,10) = 4.819$, $p< .06$, thus demonstrating less of an SOA effect than

the random object stream. One-tailed t-tests were conducted for the first four

---

[4] It was not possible to directly compare results from Chapter 2 with those from this chapter for two reasons. First, in Chapter 2 the first serial position of T2 was the T1+1 position, whereas in Experiment 8 this position was always occupied by the texture mask. Second, the experiments in each chapter used different presentation rates, and therefore, different SOA's. Consequently, the longest SOA in Chapter 2 was 534 ms (just outside the AB range) whereas the longest SOA in Chapter 5 was 853 ms (well outside the AB period). For these reasons the first SOA for Chapter 2 experiments, and the last three SOA's from all experiments was omitted from these stream by stream analyses across experiments.

SOA's in order to determine whether the mean AB magnitude was significant a given SOA. These analyses revealed that the 427 ms SOA was significant at the .05 level, $t(10) = 1.904$, $p < .05$, whilst the other three SOA's approached significance at the same level, 213 ms, $t(10) = 1.574$, $p > .05$; 320 ms, $t(10) = 1.681$, $p > .05$; 533 ms, $t(10) = 1.637$, $p > .05$. A great deal of variability was observed in the scrambled morph group and this, coupled with the relatively small AB magnitudes, accounts for the results of the analyses for this group. However, there was evidence that an AB effect was present and that it endured until 640 ms.

However, a comparison of the outcomes for the scrambled morph stream in this experiment and Experiment 1 reveal the pattern of results to be qualitatively very similar. Whilst, in Experiment 1 there was a robust AB effect, such an effect was only marginally significant here, probably because of increased variability. However, an ANOVA conducted on the two sets of data shows there to be no significant difference in AB magnitude across SOA between the two groups, $F(3,51) = .385$, $p > .10$.

The Smooth morph stream

In contrast to the other two stream types the trend analysis conducted for the smooth morph was non significant, $F(1,10) = .068$, $p > .05$ suggesting that overall there was no SOA effect. However, a significant degree of AB magnitude was revealed for the 320 ms SOA, $t(10) = 3.266$, $p < .01$ but the same was not true for the 213 ms SOA, $t(10) = 0.151$, $p > .05$; the 427 ms SOA, $t(10) = 0.513$, $p > .05$; or the 533 ms SOA, $t(10) = 0.220$, $p > .05$. Based on these results it is claimed that AB

attenuation was observed in the smooth morph condition with the exception of the 320 ms SOA.

AB attenuation was also observed in the original full stream morph experiment (Experiment 1). An ANOVA, shows there to be no difference in AB magnitude across SOA between Experiment 1 and the present experiment, $F(3,51) = 1.952$, p >.10. A qualitative analysis of both experiments reveals a similar increase in AB magnitude for the second serial position of T2 (although these correspond to different SOA's) though this effect is smaller in Experiment 1. The results of both experiments demonstrate that AB attenuation is observed when the RSVP stream uses a single object file throughout. Furthermore, this effect cannot be due to inadequate visual perceptual masking since a textured mask was used in Experiment 8.

T1 performance and the texture mask

Consideration will now be given to any differences in visual perceptual masking efficiency between this experiment, where a textured mask was used, and the original full stream experiments from Chapter 2, where the texture targets were masked by a non-textured item.

T1 performance for the smooth and scrambled morph full streams decreased to 79% and 80% respectively when the target was masked by a textured item (in Experiment 8) compared to when it was masked by a non-textured pipepan (in Experiment 1) where T1 performance was at 97% and 89% respectively. This suggests that the textured pipepan mask did serve as a more effective mask for a

pipepan target in streams comprised entirely of pipepan items, whether they were presented sequentially or not. This is supported by the fact that T1 performance was lower in the current experiment despite using a far slower presentation rate, although this reduced T1 accuracy may have resulted from resource limiting factors rather than encoding difficulty related factors (McLaughlin et al., 2001).

For the random object full stream, where T1 was masked by a non-textured random object (Experiment 2), T1 performance was at 87%. Where the same target was masked by a textured pipepan (Experiment 8), T1 performance was seen to decrease slightly to 81%. This suggests that in both experiments, despite differences in the physical nature of the mask, visual perceptual masking was about the same.

These differences in masking can be explained if one considers the perceptual nature of the streams. For both morph groups detection of the textured target involves searching among many similar but untextured pipepan items. Adding a texture to the mask makes the task more difficult because the featural nature of the target – the dotted texture – is no longer distinct. One might therefore anticipate a reduction in target detection. Now let us apply the same rationale to the random object stream. Again, when the mask is a non-textured random object, the target is featurally distinct on two counts. Firstly, as for the morph streams, it is the only item in the stream comprised of dots. Secondly, there are only two pipepans in the stream (which may constitute more of a pop-out search) and both of them are targets. In the random object stream, viewers know that if the item is a pipepan then it is relevant to the task. However, when a textured pipepan mask is used then this strategy is less useful and one might anticipate some reduction in target accuracy.

In Experiment 8 the same combination of target and mask was used for the random object and scrambled morph stream. Therefore, if it is the four critical items that determine the magnitude and shape of the AB then identical outcomes should have been observed for these two stream types. Whilst T1 performance was about the same the magnitude and shape of the AB functions were not. This outcome lends support to the notion that other items in the stream are instrumental in the size and shape of observed AB effects.

Concluding comments

The outcomes of Experiment 8 are in agreement with an intuitive analysis of Figure 23 and the experimental predictions. The random object stream, in which it is proposed that new object files are generated, led to a robust SOA effect. The scrambled object stream, in which it is supposed that object files are updated, did result in lower AB magnitudes and statistical analyses were only marginally significant. On the other hand, the smooth morph condition, where it is proposed that a single object file is maintained throughout, led to almost a complete attenuation of the AB. These results mirror those reported in Chapter 2 and confirm that those results were not due to differences in visual perceptual masking but were due to object file related factors.

Chapter 6 – General Discussion

This thesis set out to investigate the relationship between object file continuity and the AB. In the full stream experiments I manipulated the degree of object file continuity such that targets in an RSVP stream could or could not be linked to the same object file. The part stream experiments were designed to identify what aspect of the full stream paradigm modulates object file continuity and the AB. In this final chapter I will attempt to synthesise the results of these experiments. The chapter begins with a brief outline of all eight experiments presented in Chapters 2 to 5. The main outcomes arising from these experiments are then presented and discussed in relation to the concept of object files as well as other related literature. The chapter concludes with a consideration of the limitations of the thesis and suggestions for future work in this area.

Precis of thesis experimental work

Chapter 2

The original full stream RSVP experiments reported in Chapter 2 manipulated object file continuity or discontinuity such that T1 (a pipepan image comprised of big or small dots) and T2 (a pipepan image comprised of big or small squares) could or could not share the same object file.

- **Experiment 1** used a morph movie which was either presented sequentially in RSVP (smooth morph) or randomly (scrambled morph). In the smooth morph condition object file continuity was preserved and the AB was attenuated. In the scrambled condition a short-lived AB effect (recovery occurred by 267 ms) was obtained reflecting updating of object files rather than new object file generation.

- **Experiment 2** used a full stream of random objects and yielded a robust AB effect that did not recover until 467 ms reflecting the generation of new object files.

Chapter 3

Variations of the critical items paradigm were used in Chapter 3 in order to investigate factors that modulated the AB outcomes in the full stream smooth morph (where there was no AB) and random object (where there was a robust AB) conditions. A pipepan image was used to mask targets in these part stream experiments.

- **Experiment 3** utilised a pure critical items paradigm (comprising of the pipepan targets and their respective pipepan masks) and AB attenuation was again observed, suggesting that in the full morph stream the process of morphing was not a necessary condition to provide object file continuity. However, the lower level of variability in the full stream morph condition suggests that the morphing process did provide some degree of object file continuity.

- **Experiment 4** examined the role of the items intervening T1 and T2 for the random object stream. Here the four pipepan critical items paradigm was adapted by presenting a variable number of random objects between T1's mask and T2. No overall AB was obtained.

- **Experiment 5** considered the role of the items preceding T1 for both random object and smooth morph streams. Neither stream yielded an AB effect overall.

Chapter 4

In Chapter 4 the two, part stream experiments for the random object stream were repeated using a random object mask (in order to match the original full stream of random objects).

- **Experiment 6** utilised streams comprised of the two pipepan targets, random object masks and a variable number of random objects between T1's mask and T2. Despite high levels of variability a significant AB was obtained overall, suggesting that the items intervening T1 and T2 were a factor in object file continuity and AB modulation in the original full stream of random objects. An alternative explanation is that the outcome in both cases (full stream and Exp. 6) was due to using a random object mask.

- **Experiment 7** re-examined the role of the random objects that appear before T1 in the full stream condition and utilised a modified critical items paradigm comprising of the two pipepan targets, random object masks, and eleven random objects preceding T1. Contrary to theoretical predictions no AB effect was observed. This suggests that items preceding T1 played little, if any, role in AB modulation/object file continuity in the full stream of random objects (Experiment 2). This outcome also demonstrates that AB modulation in Experiment 2 was not due to the use of a random object mask.

Chapter 5

In Chapter 5 a single full stream experiment examined the efficacy and equivalence of visual perceptual masking.

- **Experiment 8** utilised full stream smooth morph, scrambled morph and random object conditions. A composite textured mask was used to ensure mask efficacy (demonstrated by the fact that T1 performance was reduced, despite the fact that the presentation rate was reduced to 9 items per second, compared to 15 items per second in Chapter 2). Trend analyses revealed AB attenuation in the smooth morph condition, a marginally significant AB in the scrambled condition and a significant AB in the random object condition. These results replicate the original full stream experiments (Experiment 1 and 2) and demonstrate that these outcomes are not due to differences in masking.

The main experimental outcomes and their implications

There is an inverse linear relationship between the AB and object file continuity

The results of the full stream RSVP experiments demonstrate there to be a link between object file continuity and modulation of the AB. That is, there is an inverse linear relationship between AB magnitude and the degree of object continuity inherent in the stream. I have shown that when RSVP targets can be linked to a common object file through the process of morphing then the AB is largely attenuated. Conversely, when

object file continuity has been partially disrupted (as in the scrambled morph full stream) but the object file can be updated and used, rather than discarded and a new one generated, less AB attenuation is obtained compared to the smooth morph condition where object continuity remains intact. At the other extreme, the full stream of random objects served to violate object file continuity to the degree that targets could not share the same object file. Consequently, the opening of a new object file (which involves attentionally demanding processes) corresponded with a loss of perceptual continuity and a robust AB effect. Rather than propose three separate categories of object file continuity (same file, update file, new file) it seems better to assume a continuum whereby the degree of object file continuity relates to AB magnitude.

Full stream outcomes are not due to differences in perceptual masking.

The differences in AB magnitude across the three types of stream are not due to differences in visual perceptual masking. This is supported by Experiment 8 in which the results of the original full stream experiments were replicated despite equating visual perceptual masking across the three stream types. Additionally, the scrambled and random object streams in the same experiment had identical combinations of targets and masks yet resulted in different outcomes. This suggests that the form and function of the AB is not solely determined by the critical items but is more likely due to object file related factors.

Having revealed a relationship between the AB and object file continuity, it was important to determine what these factors were, that is, what aspects of the full streams

facilitated object file continuity and AB modulation. Three possible candidate explanations were considered: the process of morphing, the items preceding T1 (as suggested by Raymond and Sorensen, 1994) or the items intervening T1 and T2 (which concurs with Kahneman et al.'s (1992) notion of 'recency' in object file reviewing). All three candidate explanations could account for the results of the full stream experiments. Consequently, part stream experiments were crucial in investigating these alternatives. The major outcomes from these experiments are now considered along with relevant theoretical implications.

The process of morphing was not crucial for AB attenuation

In this thesis, the process of morphing - the presentation of a complete sequence of pictures depicting one object smoothly changing into another - was found not to be a necessary condition in determining object file continuity (although its presence was associated with less variability and greater perceptual continuity). Kahneman et al.'s (1992) notion of recency in object file reviewing states that during the process of object file reviewing, the system only has access to the most recent object information (i.e. that pertaining to the visual item immediately preceding the current one). Hence, when one removes the process of morphing what remains is information that is entirely consistent with the concept of recency. This is consistent with object file continuity being preserved when the four pipepan critical items were presented alone (as in Experiment 3). The finding that the morphing process was not critical in providing object file continuity suggests that Raymond and Sorensen's (1994) use of motion may not have been crucial either. If so, then it should be possible to adopt a critical items version of

Raymond and Sorensen's experiment (trident targets and trident masks) and obtain object file continuity and AB attenuation.

Pre-T1 items are not the locus of AB modulation

Raymond and Sorensen (1994) claimed that the locus of object file continuity is the items preceding T1. In their series of trident experiments they proposed that the first distractor trident in the stream generates an object file and this file is subsequently adopted by T1 when it is also the same object (a trident). However, when T1 is a different object this necessitates changes to be made to the object file. Because the process of opening a new object file is attentionally demanding this will have a detrimental effect on subsequent object file operations (such as those required in order to process T2).

In this thesis, evidence arising from the part stream, pre-T1 experiments (Experiment 5 and 7) does not support Raymond and Sorensen's (1994) theoretical interpretation. Such an account would predict an AB effect in the random object stream. However, in Experiment 5 neither the morph nor the random object condition yielded an AB. Similarly, AB attenuation was observed when the same random object condition was presented using a random object mask (Experiment 7). These results also suggest that, at least in the random object full stream, the items preceding T1 played little role in modulating the AB effect. However, given the methodological differences between this thesis and the work of Raymond and Sorensen (e.g. they used apparent motion rather than object change), and the fact that two SOA's were significant in random object

condition of Experiment 5, the role of the pre-T1 items should not be completely eliminated.

Items intervening T1 and T2 are implicated as the locus of AB modulation

Experimental evidence arising from this thesis suggests that the items intervening T1 and T2 provide the locus for object file continuity and AB modulation. A robust AB effect was obtained when the pipepan targets (with random object masks) were presented with other random objects intervening T1 and T2 (Experiment 6). Based on this outcome one can claim that (at least for the full stream of random objects) the items intervening T1 and T2 were responsible for object file continuity and AB modulation. One can drawn upon Kahneman et al.'s (1992) concept of recency in object file reviewing in order to conceptualise the way in which the intervening items might be implicated in object file continuity. When the visual information for the pipepan T1 is overwritten by the intervening random objects then T1 and T2 can no longer share the same object file until substantial modifications have been made to it. Furthermore the degree of object file discontinuity determines the extent to which the object file needs to be modified and the resources required. These operations are consequently manifest in observed AB magnitudes.

In Experiment 4 (which presented random objects between T1 and T2 but used next consecutive pipepan items as masks) there was no AB overall, although a good sized, significant magnitude was observed at the 200 ms SOA. If, as this thesis claims, the concept of recency is implicated in object file continuity and AB modulation then one

must attempt to explain this outcome. It could be that masking a textured pipepan target with a non-textured image that is almost identical leads to processing facilitation. This may be because neither the texture nor the object information is overwritten by the mask and there is object continuity. In essence, the same object information is present for twice as long, thus enabling the object file to become consolidated (this corresponds to Kahneman et al.'s [1992] notion that object files need time to achieve specificity). These factors could have resulted in more efficient object file processes and processing of T1 which may have been complete by 267 ms thus releasing sufficient resources to accommodate T2 processing from that point onwards. In contrast, when T1 object information is overwritten by a random object (as in Experiment 6) processing of T2 will require greater resources.

Critical items (especially the +1 item) are not the whole story

Grandison et al. (1997) claim that any +1 item that serves as an effective mask will lead to an AB effect. However, the smooth morph condition in Experiment 8 is problematic for this claim because AB attenuation was observed despite the presence of an effective mask. The same condition also provides counter evidence to Shapiro et al., 1994 who argue that in order for the AB to occur, attention must be directed to an object or patterned target stimulus.

For Raymond et al. (1995) the source of interference inherent in the AB is the similarity between the +1 item and T2. Specifically, the greater the similarity between these two items the larger the AB effect. This thesis provides evidence contrary to such

a claim. The smooth morph full stream (Experiment 1 and 8) has consistently shown AB attenuation despite the presence of a very similar +1 item, and furthermore the full stream random object condition (Experiment 2) yielded a robust AB effect despite the presence of a dissimilar +1 item.

McAuliffe and Knowlton (2000) claim that AB magnitude is determined by featural complexity of the +1 item (rather than conceptual similarity or spatial area) and their work revealed a linear relationship between AB magnitude and the featural similarity of T1 and the +1 item. However, the results of Experiment 8, where different AB magnitudes were obtained across all three stream types despite no change to the featural (or conceptual) similarity of the +1 item, show that featural complexity cannot be the only factor that contributes to AB modulation.

This thesis demonstrates that AB magnitude is not simply determined by interference arising from the +1 item but that other stream items are also implicated (Isaak et al., 1999). Based on the part stream experiments the most likely candidates are the items intervening T1's mask and T2. Whilst the nature of such interference may be featural (or conceptual if we conceptualise pipes, pans and random objects as different categories) the results of the present series of experiments suggests that object file related factors are heavily implicated.

Duncan et al. (1994) established that it is possible to obtain an AB using only the critical items. However, whilst their work demonstrates that an AB can be obtained

without the remaining stream items it does not eliminate the possibility that, when present, distractor items may be a determinant of AB modulation. This thesis represents the first attempt to investigate the role of the distractors in terms of implicating the items intervening T1 and T2 in AB modulation as well as object file continuity.

Consideration of an alternative account

Visual transience

Another possible argument is that the different outcomes for the full stream experiments are due to differences in the degree of visual transience inherent in the visual displays (i.e. differences in offsets and onsets of RSVP items). In the smooth stream (where the AB was attenuated) there was virtually no visual transience in the stream compared to the random object stream (where a robust AB effect was obtained) which was characterised by a high level of visual transience. Initially, such an explanation does appear plausible. However, the same argument will not apply to the scrambled morph stream. When creating the morph it was essential to select items that were very similar in terms of size and shape in order to obtain the smooth process of change between the objects. Therefore when the morph was subsequently scrambled there was a relatively small degree of transience present across successive frames. Similarly, at least for the +1 item, the overall spatial area (and therefore visual transience) has been found to be

unimportant in determining AB magnitude (McAuliffe and Knowlton, 2000). In light of

these points an alternative explanation in terms of visual transience is not compelling.

Additionally, in the real world different objects correspond to a variety of spatial

configurations. Therefore it may be problematic, as well as and less ecologically valid,

to manipulate object file discontinuity whilst totally eliminating visual transience.

Implications for the main theories of the AB

Whilst it was not the primary aim of this thesis to provide evidence for or against

any particular theoretical account of the AB, one can consider how well this research is

accommodated within the three major AB theories – the interference model, the two-

stage model and finally, the object substitution model.

Interference model

This theoretical position holds that the four critical items (T1, T2 and their

respective masks) gain entry into VSTM and draw from a finite pool of weighting

resources in a serial manner (i.e. the resource pool diminishes as each consecutive item is

weighted). Thus, at early serial positions of T2 much of the weighting has been

consumed by T1 and the +1 item such that T2 is insufficiently weighted. Because T2 is

insufficiently weighted it becomes vulnerable to selection errors and this results in an

AB effect (Shapiro et al., 1994). Within this theory the degree of similarity between these items will result in greater interference.

Now let us consider the results of the smooth morph experiments in terms of this account. All four critical items match the object template (though they may have different textures). T1 and its mask should consume much of the resources (especially given the similarity between T1 and its mask) such that T2 is insufficiently weighted. Consequently, since these weights affect the probability that any given item is selected for report, one might anticipate decreased T2 performance when T2 is insufficiently weighted (i.e. at early serial positions).

Theoretically, an AB should have been observed but it was not. However, object file continuity inherent in the stream could have led to the critical items being processed (and perceived) as one object that changes its texture over time rather than four distinct objects that require processing individually.

In the case of the scrambled morph, all four items again match the object template but disrupted object file continuity means that each item must be weighted and processed separately. Here, the percept is no longer of one object that changes its texture over time but four slightly different objects each having different textures. In this case T2 would be insufficiently weighted and one might anticipate the AB effects which were observed for the scrambled morph stream. If competition within VSTM only involves the critical items then the same outcome should be observed for the

random object stream where the targets were masked by a pipepan (Experiment 8). The fact that this condition yielded a larger AB effect than the scrambled condition supports the claim that competition within VSTM involves more than simply the four critical items - competition may arise from distractor items following T1's mask (i.e. +2, +3, +4, etc) (Isaak et al., 1999).

Irrespective of the number of items competing within VSTM the results of Experiment 8 suggests that effects of the degree of similarity between competing items is not so clear cut. The critical items (which were identical across the scrambled and random object streams) had a high degree of similarity yet different AB profiles were observed. Similarly, even if VSTM competition involves more than four items then the scrambled morph should have resulted in greater interference than the random object stream where there is less similarity. Recent work by Kawahara, et al. (2001) also found no evidence to support the notion that a high degree of similarity between competing items results in increased interference and larger AB effects. These investigators observed substantial AB effects when targets were dissimilar.

Two-stage model

According to the two-stage model stimuli are unavailable for report until they have been processed by a second, serial stage of processing. By this account the AB occurs because T1 occupies the second stage such that T2 cannot gain access until such time that processing of T1 is complete. Similarly, the longer it takes to process T1 then the longer T2 is denied access to the second stage and it becomes more likely T2 will decay

(Chun and Potter, 1995). According to this model, any T1 target that is adequately masked should delay T1 processing in stage 2 and result in an AB. In this thesis the smooth morph yielded no AB effect despite the targets being adequately masked. This outcome is problematic for the two-stage model. Additionally, this account cannot explain the difference between the scrambled and random object outcomes where the targets and masks were identical (Experiment 8).

Object substitution

Chun and Potter's (1995) two-stage model is adapted in the object substitution model account of the AB which is centred on the role of T1 and T2 masks and the notion of object substitution (Enns and Di Lollo, 1997; Seiffert et al., 1997). By this account, while T1 occupies the second serial stage of processing T2 is vulnerable to substitution by the interruption mask that immediately follows T2.

The results of the full stream experiments here can be partially accommodated by this theoretical account. In the full stream experiments reported in Chapter 2, T2 may have been substituted by the random object mask that followed the T2 or only partially substituted by a scrambled stream mask. In the case of the smooth morph, T2 may have avoided substitution because object file continuity led to the percept of the same object changing (i.e. T2 changing into its mask) rather than the appearance of a new object. However, the same rationale could only be applied to Experiment 8 (where the textured mask was used) if we assume that it is object information that is substituted and not texture information.

Experiment 7 of this thesis presented pipepan targets, masked by random objects and eleven random objects preceding T1. Despite the fact that T1 and T2 were masked by interruption, no AB was observed. This outcome is problematic for object substitution theory.

A return to the central tenets

The AB attenuation in the smooth morph condition has implications for two of Shapiro, Arnell et al.'s (1997) central tenets identified as being common to the major AB theoretical positions. The first tenet was that greater resources are required to process T1 to the level of report when it is adequately masked. The second held that because attentional resources are finite, increased processing requirement for T1 will result in fewer resources being available in order to process T2. This in turn leads to the decay or substitution of T2 because it has not been able to attain a durable representation. On the basis of the smooth morph results presented here perhaps one might offer a caveat to these tenets: *When object file continuity is present, both target tasks may be processed as part of the same attentional episode where one object is seen to change from one state to another rather than one object being replaced by another.* Consequently, T2 does not require its own resources but capitalises on the resources allocated to T1.

Limitations of this research and future directions

This thesis has examined the relationship between the AB and object files and has also attempted to isolate which aspects of the stream modulate the AB and object file continuity. A natural progression from this thesis would be to further examine the boundaries of object file continuity using the AB paradigm. For example, if we insert scrambled morph items, or random objects between T1 and T2 do we obtain an AB effect? Additionally, what is the effect of selecting stimuli from the same versus different conceptual categories, and if we contrive object file continuity in a traditional style letter RSVP stream is the AB attenuated? Before elaborating on two of these planned experiments I would like to make a short diversion in order to consider the nature of the task used in this thesis along with the consequences of using such a task.

Throughout the series of experiments presented here the target task was to identify the size of a texture - a big or small dotted texture in the case of T1 and a big or small squared texture in the case of T2. It is notable that the defining characteristic was not intrinsic to the object (that is, pipes and pans are not generally comprised entirely of dots or squares). To avoid this criticism the majority of future work in this area should attempt to utilise target tasks involving intrinsic aspects of the stimuli such as orientation or a change in some feature of the object. Similarly, subsequent work should avoid the use of two-alternative forced-choice tasks because this leads to inflated target accuracy - T2 rarely diminishes below fifty per cent (i.e. chance) whereas it is not

uncommon for much lower T2 performance to occur using a more traditional AB paradigm. Additionally, using a measure that takes individual response biases into account, such as d′ (Green and Swets, 1966) for example, may be more appropriate.

One avenue of future research extrapolates object file research findings to the AB paradigm. Specifically, examining Gordon and Irwin's (2000) claim that within an object file, object identity is represented in an abstract form. To recapitulate, Gordon and Irwin utilised a picture preview display, followed by a word in the target display that required a lexical judgement. They found that, despite changes in the object's physical form object identity was preserved (object specific preview benefits were observed). An RSVP adaptation of this experiment would involve creating a picture RSVP stream with a lexical decision (word/non-word) T1. Where the target word (e.g. "FLOWER") relates to the preceding RSVP pictures (e.g. pictures of different flowers) one would anticipate no AB effect because object identity should be preserved because the same object file should be utilised. However, when the target word is inconsistent with the object identity held by the stream pictures this would constitute a mismatch in object identity, resulting in the generation of a new object file. This latter scenario should theoretically result in an AB effect.

Another future research idea involves manipulating object file continuity and discontinuity using streams of letters in RSVP (the traditional choice of RSVP stimuli). In the object continuity condition one would need to anchor some aspect of each stream letter such that either the same object file is used throughout or need only be updated

(rather than generate new object files). One way to achieve this would be to utilise all twelve letters of the alphabet that have an upright line of the same size (i.e. letters B, D, E, F, H, K, L, N, P, R and T). The upright bar would appear alone at the start of the stream in order to establish the object file representation. Subsequently, when the other letters appear, the perception should be of letters morphing from one to another rather than the disappearance of one letter and the reappearance of another, different, letter. Traditional target tasks could be used here, for example, identify the colour of T1 and detect the presence/absence of T2. If object file continuity can be preserved in this way then one should observe some degree of AB attenuation. In contrast, when the same target tasks are embedded in a traditional RSVP letter stream, an AB effect should be observed.

Appendix A

Experiment 1 – supplementary analysis

Mean and medium morph movie frame numbers presented before and after T1 and after T2 in the scrambled stream. Brackets denote standard deviations.

Serial position of T2 relative to T1 (SOA)	T1 position	Morph frame presented at the T1-1 position	Morph frame presented at the T+1 position	T2 position	Morph frame presented at the T2+1 position
	Fixed	Variable	Variable	Fixed	Variable
T2 position 1 (SOA 67 ms)	12			13	
mean		14 (7)			13 (7)
median		14			15
T2 position 2 (SOA 133 ms)	12			14	
mean		11 (7)	12 (8)		12 (8)
median		10	14		12
T2 position 3 (SOA 200 ms)	12			15	
mean		11 (7)	12 (7)		14 (7)
median		9	11		14
T2 position 4 (SOA 267 ms)	12			16	
mean		14 (8)	12 (7)		12 (8)
median		16	13		11

Appendix B

Experiment 5 – supplemental analyses

SCRAMBLED V SMOOTH T-TEST

Unpaired t-test for isi213
Grouping Variable: stream
Hypothesized Difference = 0
Row exclusion: EX8T-TESTDATA

	Mean Diff.	DF	t-Value	P-Value
smooth, scrambled	-6.818	20	-1.237	.2304

Group Info for isi213
Grouping Variable: stream
Row exclusion: EX8T-TESTDATA

	Count	Mean	Variance	Std. Dev.	Std. Err
smooth	11	.455	99.273	9.964	3.004
scrambled	11	7.273	234.818	15.324	4.620

Unpaired t-test for isi320
Grouping Variable: stream
Hypothesized Difference = 0
Row exclusion: EX8T-TESTDATA

	Mean Diff.	DF	t-Value	P-Value
smooth, scrambled	.727	20	.088	.9307

Group Info for isi320
Grouping Variable: stream
Row exclusion: EX8T-TESTDATA

	Count	Mean	Variance	Std. Dev.	Std. Err
smooth	11	12.909	171.891	13.111	3.953
scrambled	11	12.182	577.964	24.041	7.249

Unpaired t-test for isi427
Grouping Variable: stream
Hypothesized Difference = 0
Row exclusion: EX8T-TESTDATA

	Mean Diff.	DF	t-Value	P-Value
smooth, scrambled	-10.000	20	-1.374	.1847

Group Info for isi427
Grouping Variable: stream
Row exclusion: EX8T-TESTDATA

	Count	Mean	Variance	Std. Dev.	Std. Err
smooth	11	1.909	152.491	12.349	3.723
scrambled	11	11.909	430.291	20.743	6.254

Unpaired t-test for isi533
Grouping Variable: stream
Hypothesized Difference = 0
Row exclusion: EX8T-TESTDATA

	Mean Diff.	DF	t-Value	P-Value
smooth, scrambled	-5.909	20	-.776	.4471

Group Info for isi533
Grouping Variable: stream
Row exclusion: EX8T-TESTDATA

	Count	Mean	Variance	Std. Dev.	Std. Err
smooth	11	1.364	421.455	20.529	6.190
scrambled	11	7.273	217.018	14.732	4.442

Unpaired t-test for isi640
Grouping Variable: stream
Hypothesized Difference = 0
Row exclusion: EX8T-TESTDATA

	Mean Diff.	DF	t-Value	P-Value
smooth, scrambled	-.727	20	-.146	.8853

Group Info for isi640
Grouping Variable: stream
Row exclusion: EX8T-TESTDATA

	Count	Mean	Variance	Std. Dev.	Std. Err
smooth	11	2.636	133.255	11.544	3.481
scrambled	11	3.364	139.455	11.809	3.561

Unpaired t-test for isi746
Grouping Variable: stream
Hypothesized Difference = 0
Row exclusion: EX8T-TESTDATA

	Mean Diff.	DF	t-Value	P-Value
smooth, scrambled	.273	20	.039	.9691

Group Info for isi746
Grouping Variable: stream
Row exclusion: EX8T-TESTDATA

	Count	Mean	Variance	Std. Dev.	Std. Err
smooth	11	3.909	308.291	17.558	5.294
scrambled	11	3.636	222.855	14.928	4.501

Unpaired t-test for isi853
Grouping Variable: stream
Hypothesized Difference = 0
Row exclusion: EX8T-TESTDATA

	Mean Diff.	DF	t-Value	P-Value
smooth, scrambled	8.000	20	1.261	.2218

Group Info for isi853
Grouping Variable: stream
Row exclusion: EX8T-TESTDATA

	Count	Mean	Variance	Std. Dev.	Std. Err
smooth	11	5.091	171.891	13.111	3.953
scrambled	11	-2.909	270.891	16.459	4.963

SCRAMBLED V RANDOM OBJECTS – T-TESTS

Unpaired t-test for isi213
Grouping Variable: stream
Hypothesized Difference = 0
Row exclusion: EX8T-TESTDATA

	Mean Diff.	DF	t-Value	P-Value
scrambled, random	-8.000	20	-1.346	.1934

Group Info for isi213
Grouping Variable: stream
Row exclusion: EX8T-TESTDATA

	Count	Mean	Variance	Std. Dev.	Std. Err
scrambled	11	7.273	234.818	15.324	4.620
random	11	15.273	153.818	12.402	3.739

Unpaired t-test for isi320
Grouping Variable: stream
Hypothesized Difference = 0
Row exclusion: EX8T-TESTDATA

	Mean Diff.	DF	t-Value	P-Value
scrambled, random	-12.455	20	-1.414	.1728

Group Info for isi320
Grouping Variable: stream
Row exclusion: EX8T-TESTDATA

	Count	Mean	Variance	Std. Dev.	Std. Err
scrambled	11	12.182	577.964	24.041	7.249
random	11	24.636	275.655	16.603	5.006

Unpaired t-test for isi427
Grouping Variable: stream
Hypothesized Difference = 0
Row exclusion: EX8T-TESTDATA

	Mean Diff.	DF	t-Value	P-Value
scrambled, random	4.909	20	.576	.5707

Group Info for isi427
Grouping Variable: stream
Row exclusion: EX8T-TESTDATA

	Count	Mean	Variance	Std. Dev.	Std. Err
scrambled	11	11.909	430.291	20.743	6.254
random	11	7.000	367.400	19.168	5.779

Unpaired t-test for isi533
Grouping Variable: stream
Hypothesized Difference = 0
Row exclusion: EX8T-TESTDATA

	Mean Diff.	DF	t-Value	P-Value
scrambled, random	-7.091	20	-1.160	.2598

Group Info for isi533
Grouping Variable: stream
Row exclusion: EX8T-TESTDATA

	Count	Mean	Variance	Std. Dev.	Std. Err
scrambled	11	7.273	217.018	14.732	4.442
random	11	14.364	194.255	13.938	4.202

Unpaired t-test for isi640
Grouping Variable: stream
Hypothesized Difference = 0
Row exclusion: EX8T-TESTDATA

	Mean Diff.	DF	t-Value	P-Value
scrambled, random	4.000	20	.700	.4918

Group Info for isi640
Grouping Variable: stream
Row exclusion: EX8T-TESTDATA

	Count	Mean	Variance	Std. Dev.	Std. Err
scrambled	11	3.364	139.455	11.809	3.561
random	11	-.636	219.455	14.814	4.467

Unpaired t-test for isi746
Grouping Variable: stream
Hypothesized Difference = 0
Row exclusion: EX8T-TESTDATA

	Mean Diff.	DF	t-Value	P-Value
scrambled, random	2.455	20	.466	.6459

Group Info for isi746
Grouping Variable: stream
Row exclusion: EX8T-TESTDATA

	Count	Mean	Variance	Std. Dev.	Std. Err
scrambled	11	3.636	222.855	14.928	4.501
random	11	1.182	81.764	9.042	2.726

Unpaired t-test for isi853
Grouping Variable: stream
Hypothesized Difference = 0
Row exclusion: EX8T-TESTDATA

	Mean Diff.	DF	t-Value	P-Value
scrambled, random	-6.273	20	-.894	.3822

Group Info for isi853
Grouping Variable: stream
Row exclusion: EX8T-TESTDATA

	Count	Mean	Variance	Std. Dev.	Std. Err
scrambled	11	-2.909	270.891	16.459	4.963
random	11	3.364	271.255	16.470	4.966

SMOOTH V RANDOM OBJECTS – T-TESTS

Unpaired t-test for isi213
Grouping Variable: stream
Hypothesized Difference = 0
Row exclusion: EX8T-TESTDATA

	Mean Diff.	DF	t-Value	P-Value
smooth, random	-14.818	20	-3.089	.0058

Group Info for isi213
Grouping Variable: stream
Row exclusion: EX8T-TESTDATA

	Count	Mean	Variance	Std. Dev.	Std. Err
smooth	11	.455	99.273	9.964	3.004
random	11	15.273	153.818	12.402	3.739

Unpaired t-test for isi320
Grouping Variable: stream
Hypothesized Difference = 0
Row exclusion: EX8T-TESTDATA

	Mean Diff.	DF	t-Value	P-Value
smooth, random	-11.727	20	-1.839	.0809

Group Info for isi320
Grouping Variable: stream
Row exclusion: EX8T-TESTDATA

	Count	Mean	Variance	Std. Dev.	Std. Err
smooth	11	12.909	171.891	13.111	3.953
random	11	24.636	275.655	16.603	5.006

Unpaired t-test for isi427
Grouping Variable: stream
Hypothesized Difference = 0
Row exclusion: EX8T-TESTDATA

	Mean Diff.	DF	t-Value	P-Value
smooth, random	-5.091	20	-.741	.4676

Group Info for isi427
Grouping Variable: stream
Row exclusion: EX8T-TESTDATA

	Count	Mean	Variance	Std. Dev.	Std. Err
smooth	11	1.909	152.491	12.349	3.723
random	11	7.000	367.400	19.168	5.779

Unpaired t-test for isi533
Grouping Variable: stream
Hypothesized Difference = 0
Row exclusion: EX8T-TESTDATA

	Mean Diff.	DF	t-Value	P-Value
smooth, random	-13.000	20	-1.738	.0977

Group Info for isi533
Grouping Variable: stream
Row exclusion: EX8T-TESTDATA

	Count	Mean	Variance	Std. Dev.	Std. Err
smooth	11	1.364	421.455	20.529	6.190
random	11	14.364	194.255	13.938	4.202

Unpaired t-test for isi640
Grouping Variable: stream
Hypothesized Difference = 0
Row exclusion: EX8T-TESTDATA

	Mean Diff.	DF	t-Value	P-Value
smooth, random	3.273	20	.578	.5697

Group Info for isi640
Grouping Variable: stream
Row exclusion: EX8T-TESTDATA

	Count	Mean	Variance	Std. Dev.	Std. Err
smooth	11	2.636	133.255	11.544	3.481
random	11	-.636	219.455	14.814	4.467

Unpaired t-test for isi746
Grouping Variable: stream
Hypothesized Difference = 0
Row exclusion: EX8T-TESTDATA

	Mean Diff.	DF	t-Value	P-Value
smooth, random	2.727	20	.458	.6519

Group Info for isi746
Grouping Variable: stream
Row exclusion: EX8T-TESTDATA

	Count	Mean	Variance	Std. Dev.	Std. Err
smooth	11	3.909	308.291	17.558	5.294
random	11	1.182	81.764	9.042	2.726

Unpaired t-test for isi853
Grouping Variable: stream
Hypothesized Difference = 0
Row exclusion: EX8T-TESTDATA

	Mean Diff.	DF	t-Value	P-Value
smooth, random	1.727	20	.272	.7883

Group Info for isi853
Grouping Variable: stream
Row exclusion: EX8T-TESTDATA

	Count	Mean	Variance	Std. Dev.	Std. Err
smooth	11	5.091	171.891	13.111	3.953
random	11	3.364	271.255	16.470	4.966

References

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

Allport, A., Styles, E. A., & Hseih, S. (1994). Shifting intentional set: Exploring the dynamic control of tasks. In C. Umilta & M. Moscovitch (Eds.), *Attention and Performance XV: Conscious and nonconscious information processing* (pp. 421-452). Cambridge, MA: MIT Press.

Baylis, G. C., & Driver, J. (1993). Visual attention and objects: Evidence for heirarchical coding of location. *Journal of Experimental Psychology: Human Perception and Performance, 19*(3), 451-470.

Bornemann, E. (1942). Untersuchungen uber den Grad der geistigen Beanspruchung. *Zeitschrift fur Arbeitshysiologie*, 142-172.

Brehaut, J.C., Enns, J. T., & Di Lollo, V. (1999) Visual masking plays two roles in the attentional blink. *Perception and Psychophysics, 61(7),* 1436-1448

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

Chelazzi, L., Miller, E. K., Duncan, J., & Desimone, R. (1993). Dual mechanisms of short-term memory: Ventral prefrontal cortex. *Society for Neuroscience Abstracts, 19*, 975.

Chun, M. (1997). Temporal binding errors are redistributed by the attentional blink. *Perception and Psychophysics, 59(8)*, 1191-1199

Chun, M. M., Bromberg, H. S., & Potter, M. C. (1994). *Conceptual similarity between targets and distractors in the attentional blink.* Paper presented at the 35th annual meeting of the Psychonomic Soceity, St. Louis.

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*(1), 109-127.

Creamer, L. R. (1963). Event uncertainty, psychological refractory period, and human data processing. *Journal of Experimental Psychology, 66,* 187-194.

Davis, G., Driver, J., Pavani, F., & Shepherd, A. (2000). Reappraising the apparent costs of attending to two separate visual objects. *Vision Research, 40,* 1323-1332.

Duncan, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology: General, 113,* 501-517.

Duncan, J. (1996). Cooperating brain systems in slective perception and attention. In T. Inui & J. L. McClelland (Eds.), *Attention and Performance XVI* (pp. 549-578). London: MIT Press.

Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review, 96*(3), 433-458.

Duncan, J., & Humphreys, G. W. (1992). Beyond the search surface: Visual search and attentional engagement. Journal of Experimental Psychology: Human Perception and Performance, 18, 578-588.

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

Egly, R., Driver, J., & Rafal, R. D. (1994). Shifting visual attention between objects and locations: evidence from normal and parietal lesion subjects. *Journal of Experimental Psychology: General, 123*, 167-177.

Enns, J. T., Brehaut, J., & Shore, D. I. (1996). Attended objects are on view longer than unattended ones. *Abstracts of the Psychonomic Soceity, 1*, 40.

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

Forster, K. I. (1970). Visual perception of rapidly presented word sequences of varying complexity. *Perception and Psychophisics, 8*, 215-221.

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

Gopher, D., Brickner, M., & Navon, D. (1982). Different difficulty manipulations interact differentially with task emphasis: Evidence for multiple resources. *Journal of Experimental Psychology: Human Perception and Performance, 8*, 146-157.

Gordon, R. D., & Irwin, D. E. (1996). What's in an object file? Evidence from priming studies. *Perception and Psychophysics, 58*(8), 1260-1277.

Gordon, R. D., & Irwin, D. E. (2000). The role of physical and conceptual properties in preserving object continuity. *Journal of Experimental Psychology: Learning, Memory and Cognition, 26*(1), 136-150.

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

Green, D.M., & Swets, J.A. (1966). *Signal detection theory and psychophysics.* New York: Wiley.

Henderson, J. M. (1994). Two representational systems in dynamic visual identification. *Journal of Experimental Psychology: General, 123*, 410-426.

Henderson, J. M., & Anes, M. D. (1994). Roles of object-file review and type priming in visual identification within and across eye fixations. *Journal of Experimental Psychology: Human Perception and Performance, 20*(4), 826-839.

Hommel, B. (1998). Event Files: Evidence for automatic integration of stimulus-response episodes. *Visual Cognition, 5*(1/2), 183-216.

Howell, D.C. (1997). *Statistical methods for psychology.* London: Wandsworth Publishing Company.

Irwin, D. E. (1992). Memory for position and identity across eye movements. *Journal of Experimental Psychology: Learning, Memory and Cognition, 18*, 307-317.

Irwin, D. E., & Andrews, R. V. (1996). Integration and accumulation of information across Saccadic eye movements. In T. Inui & J. L. McClelland (Eds.), *Attention and Performance XVI* (pp. 125-156). London: MIT Press.

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

Jolicoeur, P. (1998). Modulation of the attentional blink by on-line response selection: Evidence from speeded and unspeeded Task1 decisions. *Memory and Cognition, 26*(5), 1014-1032.

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

Kahneman, D., Treisman, A., & Gibbs, B. J. (1992). The reviewing of object files: Object-specific integration of information. *Cognitive Psychology, 24*, 175-219.

Kawahara, J-I., Di Lollo, V., & Enns, J. T. (2001). Attentional requirements in visual detection and identification: Evidence from the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance, 27*(4), 969-984.

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

Luck, S. J., Vogel, E. K., & Shapiro, K. L. (1996). Word meanings can be accessed but not reported during the attentional blink. *Nature, 383*(17 October), 616-618.

Martin, J., Isaak, M. I., & Shapiro, K. L. (1995). *Probe identification errors support an interference model of the attentional blink in rapid serial visual presentation.* Paper presented at the Annual meeting of the American Psychological Society, New York.

McAuliffe, S. P., & Knowlton, B. J. (2000). Dissociating the effects of featural and conceptual interference on multiple target processing in rapid serial visual presentation. *Perception and Psychophysics, 62*(1), 187-195.

McLaughlin, E. N., Shore, D. I., & Klein, R. M. (2001). The attentional blink is immune to masking-induced data limits. *The Quarterly Journal of Experimental Psychology: Human Experimental Psychology, 54A*(1), 169-196.

Moran, J., & Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science, 229*, 782-784

Mordkoff. (1999). *Divided attention and redundancy gains*, [web page]. http://gandalf.la.psu.edu/catnap/div-attn.html.

Navon, D., & Miller, J. (1987). Role of outcome conflict in dual-task interference. *Journal of Experimental Psychology: Human Perception and Performance, 13*(435-448), 435-448.

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

Pashler, H. (1990). Coordinate frame for symmetry detection and object recognition. *Journal of Experimental Psychology: Human Perception and Performance, 16*, 150-163.

Pashler, H. (1994). Dual-task interference in simple tasks: Data and theory. *Psychological Bulletin, 116*(2), 220-244.

Pashler, H., Carrier, M., & Hoffman, J. (1993). Saccadic eye movements and dual task interference. *Quarterly Journal of Experimental Psychology, 46A*, 51-82.

Pashler, H. E. (1998). *The psychology of attention*. London: MIT Press.

Posner, M.I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology, 32*, 3-25.

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

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*(3), 849-860.

Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1995). Similarity determines the attentional blink. *Journal of Experimental Psychology, 21*(3), 653-662.

Raymond, J. E., & Sorensen, R. E. (1994). *Motion, object selection, and the attentional blink.* Paper presented at the 35th annual meeting of the Psychonomic Society, St. Louis.

Reeves, A., & Sperling, G. (1986). Attention gating in short-term visual memory. *Psychological Review, 94*, 180-206.

Rensink, R.A., O'Regan, J.K., & Clark, J.J. (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science, 8,* 368-373

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

Seiffert, A. E., & Di Lollo, V. (1997). Low-level masking in the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance, 23*(4), 1061-1073.

Shapiro, K. L., Arnell, K. M., & Raymond, J. E. (1997). The attentional blink. *Trends in Cognitive Sciences, 1*(8), 291-296.

Shapiro, K. L., Caldwell, J., & Sorensen, R. E. (1997). Personal names and the attentional blink: A visual "Cocktail Party" effect. *Journal of Experimental Psychology: Human Perception and Performance, 23*(2), 504-514.

Shapiro, K., Driver, J., Ward, R., & Sorensen, R. E. (1997). Priming from the attentional blink: A failure to extract visual tokens but not visual types. *Psychological Science, 8*(2), 95-100.

Shapiro, K. L., & Raymond, J. E. (1994). Temporal allocation of visual attention. In D. Dagenbach & T. Carr (Eds.), *Inhibitory processes in attention, memory and language* (pp. 151-188). New York: Academic Press.

Shapiro, K. L., Raymond, J. E., & Arnell, K. M. (1994). Attention to visual pattern information produces the attentional blink in rapid serial visual presentation. *Journal of Experimental Psychology: Human Perception and Performance, 20*(2), 357-371.

Shepard, R. (1984). Ecological constraints on internal representations: Resonant kinematics of perceiving, imagining and dreaming. *Psychological Review, 91*, 417-447.

Sheppard, D.M., Duncan, J., Shapiro, K.L., & Hillstrom, A.P. (in press). Objects and events in the attentional blink. *Psychological Science.*

Shih, H-I. (2000). Recall of two visual targets embedded in RSVP streams of distractors depends on their temporal and spatial relationship. *Perception and Psychophysics, 62(7),* 1348-1355.

Smith, S. L. (1962). Color coding and visual search. *Journal of Experimental Psychology, 64*, 434-440.

Snodgrass, J. G., & Vanderwart, M. (1980). A standardized set of 260 pictures: Norms for name agreement, image agreement, familiarity, and visual complexity. *Journal of Experimental Psychology: Human Learning, and Memory, 6*, 174-215.

Spelke, E., Hirst, W., & Neisser, U. (1976). Skills of divided attention. *Cognition, 4*, 215-230.

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs: General and Applied, 74*(498), 1-29.

Sperling, G., & Melchner, M. J. (1978). Visual search, visual attention, and the attentional operating characteristics. In J. Requin (Ed.), *Attention and Performance VII* (pp. 675-686). Hillsdale, NJ: Erlbaum.

Sperling, G., & Weichselgartner, E. (1995). Episodic theory of the dynamics of spatial attention. *Psychological Review, 102*, 503-532.

Styles, E. (1997). *The psychology of attention*. Hove, East Sussex: Psychological Press Limited.

Taylor, T. L., & Hamm, J. (1997). Category effects in temporal visual search. *Canadian Journal of Experimental Psychology, 30*, 79-109.

Telford, C. W. (1931). The refractory phase of voluntary and associative responses. *Journal of Experimental Psychology, 14*, 1-36.

Ternus, J. (1938). The problem of phenomenal identity. In W., D. Ellis (Ed.)*A source book of Gestalt psychology*. New York: Harcourt, Brace and Co.

Treisman, A. (1986a). Features and objects in visual processing. *Scientific American, 255*, 114-125.

Treisman, A. (1986b). Properties, parts and objects. In K. R. Boff, L. Kaufman, & J. P. Thomas (Eds.), *Handbook of Perception and Human Performance* (pp. 35-70). New York: Wiley.

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

Tsal, Y., Meiran, N., & Lamy, D. (1995) Towards a resolution theory of visual attention. *Visual Cognition, 2*, 313-330

Ullman, S. (1979). *The interpretation of visual motion*. Cambridge, MA: MIT Press.

Visser, T.A.W., Bischof, W.F., & Di Lollo, V. (1999) Attentional switching in spatial and nonspatial domains: Evidence from the Attentional Blink. *Psychological Bulletin, 125(4)* 458-469

Vogel, E.K., Luck, S.J., & Shapiro, K.L. (1998). Electrophysiological evidence for a postperceptual locus of suppression during the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance, 24(6),* 1656-1674.

Watson, S. E., & Kramer, A. F. (1999). Object-based visual selective attention and perceptual organization. *Perception and Psychophysics, 61*(1), 31-49.

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. *British Journal of Psychology, 43*, 2-19.

Welford, A. T. (1967). Single-channel operation in the brain. *Acta Psychologica, 27*, 5-22.

Wickens, C. D. (1984). Processing resources in attention, dual task performance, and workload assessment. In R. Parasuraman & R. Davies (Eds.), *Varieties of Attention* (pp. 63-102). New York: Academic Press.

Wickens, C. D., Sandry, D., & Vidulich, M. (1983). Compatability and resource competition between modalities of input, central processing, and output: Testing a model of complex performance. *Human Factors, 25*, 227-248.