



ELSEVIER

Contents lists available at ScienceDirect

Cognitive Psychology

journal homepage: www.elsevier.com/locate/cogpsych

Too much control can hurt: A threaded cognition model of the attentional blink

Niels A. Taatgen^{a,b,*}, Ion Juvina^a, Marc Schipper^{c,d,1}, Jelmer P. Borst^c, Sander Martens^{c,d}

^aDepartment of Psychology, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15123, United States

^bDepartment of Artificial Intelligence, University of Groningen, PO Box 72, 9700 AB Groningen, The Netherlands

^cNeuroImaging Center, University of Groningen, PO Box 72, 9700 AB Groningen, The Netherlands

^dUniversity Medical Center Groningen, Hanzeplein 1, PO Box 11120, 9700 CC Groningen, The Netherlands

ARTICLE INFO

Article history:

Accepted 28 December 2008

Available online xxx

Keywords:

Attentional blink
Multi-tasking
Cognitive model
Cognitive control

ABSTRACT

Explanations for the attentional blink (AB; a deficit in identifying the second of two targets when presented 200–500 ms after the first) have recently shifted from limitations in memory consolidation to disruptions in cognitive control. With a new model based on the threaded cognition theory of multi-tasking we propose a different explanation: the AB is produced by an overexertion of control. This overexertion is produced by a production rule that blocks target detection during memory consolidation. In addition to fitting many known effects in the literature, the model predicts that adding certain secondary tasks will decrease the AB. In Experiment 1, a secondary task is added to the AB task in which participants have to respond to a moving dot. As predicted, AB decreases. Experiment 2 expands this result by controlling for learning, and adds a second variation, rotating the first target. For this variation the model predicts an increase in AB, which is indeed what we found.

© 2009 Elsevier Inc. All rights reserved.

1. Introduction

The term *cognitive control* is used to refer to cognitive processes that help us focus on our goals and plans, and prevent external stimuli and events from interfering with them. Cognitive control is also needed to coordinate multiple components in executing a task, especially when the task is new. Rapid

* Corresponding author. Address: Department of Psychology, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15123, United States. Fax: +1 412 268 2844.

E-mail address: taatgen@cmu.edu (N.A. Taatgen).

¹ Marc Schipper is currently at the Department of Human Neurobiology, University of Bremen.

Serial Visual Presentation (RSVP) tasks are challenging from the perspective of control (Raymond, Shapiro, & Arnell, 1992). In these tasks, a rapid stream of visual stimuli, typically at a rate 100 ms per stimulus, is presented to participants, whose task is to identify and report up to two targets from that stream. The RSVP task entails a combination of two basic cognitive steps: detecting visual targets that satisfy some criterion, and memorizing the targets. We may assume that both basic steps are existing skills, but that the combination of the two is novel. It is up to control mechanisms to prioritize and sequence the two steps and assign the appropriate cognitive resources to them. As such, control serves as cognitive glue to build new skills from existing skills. In the case of RSVP tasks, this is not always completely successful, as is evidenced by the fact that the first target is often identified correctly, but reporting the second is impaired if the interval between the two targets is between 200 and 500 ms, a phenomenon called the *attentional blink* (AB). A possible alternative explanation for the AB is that our cognitive system is not up to this task, and that it is not capable of handling information at such a fast rate. However, this is contradicted by *Lag-1 Sparing*: if the two targets are up to approximately 100 ms apart, the report of the second target is not impaired. Moreover, recent studies have shown that people are capable of reporting even longer sequences of three or more targets without any AB (Di Lollo, Kawahara, Ghorashi, & Enns, 2005; Nieuwenstein & Potter, 2006).

In this article we present a model that attributes the attentional blink to cognitive control. Although there are several other models based on control (Di Lollo et al., 2005; Olivers & Meeter, 2008), our model takes a new perspective in claiming that the AB is due to an overexertion of control. According to the model, the AB is caused by the activation of a production rule that suppresses target detection during memory consolidation. As we will see, this model allows us to offer an explanation why some individuals exhibit no blink at all (Feinstein, Stein, Castillo, & Paulus, 2004; Martens, Johnson, Bolle, & Borst, 2009; Martens, Munneke, Smid, & Johnson, 2006; Martens & Valchev, 2009). It can also explain why certain manipulations to decrease control in the task decrease the blink (Arend, Johnston, & Shapiro, 2006; Olivers & Nieuwenhuis, 2005, 2006). Using this, it will produce a novel prediction that adding a concurrent task decreases the blink, but manipulating another aspect of the task will increase it.

We will first give an overview of existing models, and then explain how our own model produces the basic AB. We will subsequently show how the model can explain a number of other standard findings in the AB literature, before discussing two experiments in which the RSVP task is combined with a concurrent dot-detection task, and a version in which the first target is rotated.

1.1. Models of the attentional blink

Although there are many theories and models of the AB, they share many similarities. Many models distinguish between a fast and parallel target detection process and a slow and serial memory-consolidation process, and attribute the AB to the interaction between the two. In Shapiro, Raymond, and Arnell (1994); see also Isaak, Shapiro, and Martin (1999) interference model, the memory-consolidation process is resource-limited and requires between 200 and 500 ms to consolidate a target. Lag-1 sparing is explained by the fact that both the first target (T1) and its immediate successor (T1+1) are always transferred to working memory. If T1+1 is the second target, working memory only has to deal with T1 and T1+1, but if the second target is later in the stream, working memory has to resolve competition between T1, T1+1, T2 (the second target), and T2+1 (the successor of the second target), leaving insufficient resources for T2 on the shorter lags. Chun and Potter's (1995) two-stage model has a slightly different explanation that also focuses on memory consolidation. It assumes that once consolidation of the first target has started, any subsequent targets fail to reach the consolidation stage until consolidation of the first target is done.

Although many of the original models of the AB describe it in terms of cognitive information processing, they do not have actual simulations that reproduce the data. Newer models do produce such simulations, often using neural networks or mathematical modeling. Bowman and Wyble (2007) provide an excellent overview of many of these models, and we will review a few of them here. The global workspace model by Dehaene, Sergent, and Changeux (2003) provides a detailed neural simulation of the AB phenomenon. In their model, visual stimuli have to progress through several processing stages in order to reach the global workspace from where it can be reported. However, once T1 reaches the global workspace, it inhibits subsequent visual input, suppressing T2 if it follows too soon after T1. The

model produces some “Lag-0” sparing when T1 and T2 are presented simultaneously, because in that case both targets reach the global workspace at the same time, but it fails to explain Lag-1 sparing.

Both the corollary discharge of attention movement (CODAM) model by [Fragopanagos, Kockelkoren, and Taylor \(2005\)](#) and the simultaneous type, serial token (ST²) model by [Bowman and Wyble \(2007\)](#) are neural network versions of a two-stage model. In the CODAM model, visual inputs are represented in an object map. In order for a visual input to progress to working memory, it needs a boost from an attentional system that is linked to both the current goal and progress in updating working memory. While working memory is tied up consolidating T1, the attentional system withholds its boost, causing T2 to be lost. The ST² model uses different mechanisms to implement a two-stage model. In ST², visual stimuli have to pass through a saliency filter that extracts targets from the distractors. After that, they are available as types that have to be bound to tokens in a binding pool in order to be reported later. This binding process is facilitated by a transient attentional enhancement (TAE) that lasts for approximately 100 ms, and is only available again after memory consolidation has taken place. Lag-1 sparing can be attributed to T2 falling within the TAE window. The idea of attentional enhancement is also present in a model by [Nieuwenhuis, Gilzenrat, Holmes, and Cohen \(2005\)](#), in which this function is attributed to the Locus Coeruleus (LC). Contrary to the CODAM and ST² models, the LC has a refractory period that is not directly tied to the memory-consolidation process, even though it is similar in duration.

[Shih's \(2008\)](#) cascade model is a mathematical model of the AB. In the cascade model stimuli can make it into an attentional window through two pathways: top-down selection, based on the task criteria, or bottom-up saliency. The attentional window typically has a capacity of two items (although it can vary depending on task instructions), and once it fills up, its contents are transferred to working memory. Working memory is constrained by limitations similar to other models: once it is busy consolidating the contents of a particular attentional window, subsequent content has to wait and is subject to decay.

One aspect shared by these models is that in some way memory consolidation not keeping up with target detection produces the AB. Lag-1 sparing should be a sign that this idea might have problems, but most models fix this by allowing T1 and T1+1 to be consolidated at the same time.

A further challenge to the idea of working memory as a bottleneck comes from experiments by [Di Lollo et al. \(2005\)](#); also see [Olivers, van der Stigchel, and Hulleman \(2007\)](#). In a variation of the classical AB experiment, they gave participants an RSVP string with three targets (letters, in this case) in sequence (e.g., 425293TCG2394), and required them to report all three targets, contrasting this with the standard Lag-2 blink situation (e.g., 425293T3G2394). Surprisingly enough, accuracy for reporting T1 was the same as the accuracy for reporting T3 in the three-target sequence (with T2 accuracy even higher than either), even though the standard blink sequence produced the standard blink (i.e., T2 accuracy lower than T1 accuracy). If the AB is a working memory problem, memorizing three targets should be even more of a problem than two targets separated by a distractor. [Di Lollo et al. \(2005\)](#) therefore propose a different explanation for the AB: a temporary loss of control (TLC). This account assumes that detecting targets and rejecting distractors requires active control of an input filter. Once T1 is detected, control on the input filter diminishes in order to further process T1. If T1+1 is also a target, the input filter's configuration remains unaltered, allowing the next target and potentially subsequent targets to be processed. If T1+1 is a distractor, on the other hand, control on the input filter is temporarily disrupted, leading to improper processing of subsequent items, possibly producing a blink.

[Olivers and Meeter's \(2008\)](#) Boost and Bounce model has a similar proposition. In their model, detection of T1 gives a boost to the processing of the next item, T1+1. If T1+1, however, is a distractor, the additional activation of the item produces a backlash (the “bounce”) that inhibits processing of the next few items, which leads to an AB.

Both the [Di Lollo et al. \(2005\)](#) and the [Olivers and Meeter \(2008\)](#) account shift the blame of the AB from memory consolidation to cognitive control. This fits well with our goal of explaining a dramatic decrease of AB due to individual differences or a distracting task, because cognitive control is configurable, while the more structural limitations of working memory are not. However, Di Lollo et al.'s explanation is in terms of a loss of control, which would predict that if there are other processes that also require control, less control could be committed to the RSVP task, leading to poorer performance and a potentially increased AB.

We, on the other hand, believe that there are conditions under which the addition of other tasks and the resulting decrease in control will actually lead to better performance, that is, less AB.

2. The threaded cognition model

The model we present attributes the AB to a control process that suspends target detection when there is an apparent conflict between target detection and memory consolidation. It combines the viewpoint of models that attribute the blink to control like the TLC model (Di Lollo et al., 2005) and the Boost and Bounce model (Olivers & Meeter, 2008) with the viewpoint of two-stage models that state that the blink arises out of a conflict between fast target detection and slow memory consolidation (Bowman & Wyble, 2007; Chun & Potter, 1995). Instead of Di Lollo et al.'s loss of control we attribute the blink to an overexertion of control that is exhibited when target detection and memory consolidation are combined. In our model, we see the RSVP task as a combination of two pre-existing skills: detecting targets and consolidating items in memory. Given the nature of the task, target detection and memory consolidation have to operate in parallel, and can therefore be considered as subtasks in a multi-tasking situation in which the two tasks compete for cognitive resources. To model multi-tasking, Salvucci and Taatgen (2008) have developed *threaded cognition*. Threaded cognition assumes a set of cognitive resources or modules that can all operate in parallel. However, a single resource can only be used for a single task at a time. If no additional control is exerted, each task or thread uses resources according to a greedy/polite policy: if a resource is available and the thread needs it, it will take it. As soon as the thread has no longer a need for a resource, it relinquishes the resource. The cognitive resources that play a role in the model are modules from the ACT-R architecture (Anderson, 2007):

- A visual module, to perceive the string of items (the input).
- Procedural memory, where conditions provided by the other modules are mapped onto actions. Knowledge in procedural memory is represented by production rules.
- Declarative memory, which is used to determine whether an item is a target or a distractor.
- The imaginal module, which for the purposes of this model acts as a limited working memory store, and plays the central role in memory consolidation.

Fig. 1 outlines the subtasks of target detection and memory consolidation in terms of the resources they use on a time line. In these diagrams, the resources are laid out in the rows, while the horizontal axis represents time. The exception is the top row in the diagram, which is not a resource but represents the current information on the display. A box indicates that a resource is in use at the time interval corresponding to the time axis, and arrows indicate functional dependencies between the boxes. Target detection is depicted in Fig. 1a and b. Fig. 1a shows an example of dismissing a distractor. First,

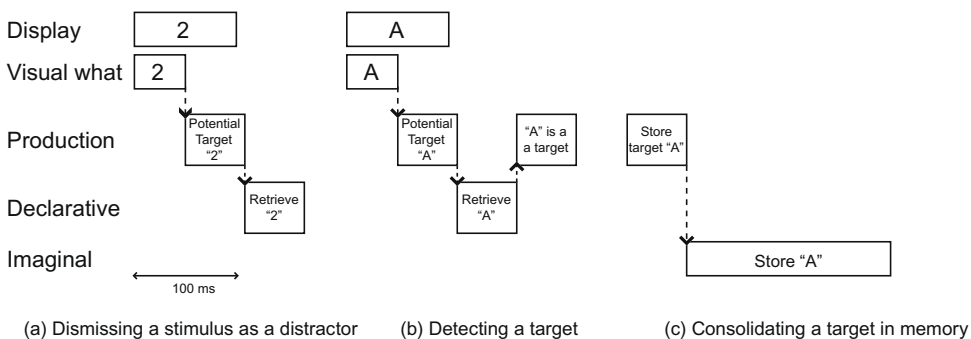


Fig. 1. Module activity chart for the two subtasks: (a) target detection on a distractor, (b) target detection on a target, and (c) memory consolidation.

the number 2 appears on the display, after which the visual resource starts processing the visual information. After the identity of the character has been determined, the procedural resource initiates a search in declarative memory to determine the category of the character. If the category of the character is “digit”, it means it is a distractor, and target detection is allowed to continue. This continuation is not achieved by a separate rule, but is incorporated in the “potential-target” rule: it has as a condition that the retrieval system is either free, or that a distractor (i.e., the previous item that is still in the retrieval system) has been retrieved. Fig. 1b depicts the situation in which a target is recognized. It is similar to Fig. 1a, except that recognition of a target triggers a production rule. Fig. 1c shows an example of memory consolidation: a production rule initiates the storage process in the imaginal module, which then starts a consolidation step of on average 250 ms.

When existing skills are combined into a new skill, it may be necessary to add some form of control to coordinate appropriate execution (Taatgen, 2005). Because the structure of the task suggests that memory consolidation follows target detection, a control production rule (the “protect consolidation” rule) has been added to the model to suppress target detection when there is an apparent conflict between target detection and memory consolidation. This apparent conflict occurs when declarative memory has finished retrieving an item while memory consolidation is in process. The production rule for resolving the conflict, however, is given a low priority (*utility* in ACT-R’s terms), so it only triggers when no production with a higher utility matches. This means that in practice it will activate during the consolidation of T1 and after the retrieval of a distractor. If a target is retrieved, the “X is a target” rule with a higher utility will activate instead. When the control rule does trigger it suspends target detection until memory consolidation has finished. Note that suspension does not imply that the input is not processed at all: the input is still processed by the visual module, and may provoke additional processing and spreading of activation that is unrelated to task. The suspension of target detection can produce an AB if T2 appears during memory consolidation. This process is somewhat similar to the Boost and Bounce model (Olivers & Meeter, 2008), but for different reasons. We call it an *apparent* conflict: target detection and memory consolidation do not share resources (apart from procedural), and can therefore operate in parallel. If the control production, for whatever reason, does not fire, the blink disappears.

The two subtasks, memory consolidation and the control production, each of which has to be carried out multiple times, are fit together like a jigsaw puzzle according to the principles of threaded cognition: if a step needs a resource and it is available, it can use the resource (greediness), but as soon as it is done, it has to relinquish it (politeness). Fig. 2 shows two examples of the model’s behavior for Lag-3 (panel a) and 1 (panel b).

Important to understanding Fig. 2 is that in the “Visual what” row a particular perceived item is available for as long as the next item has not been encoded completely. For example, the first “2” in the stream is available when the “2” on the display has been encoded (at approximately 50 ms), and is available for further processing until the end of the encoding of the “A” (at approximately 150 ms).

In the Lag-3 example in Fig. 2a, perception of the letter “A” initiates a declarative memory retrieval to verify its category. The “A” is then recognized as a target, after which memory consolidation (Store “A”) is initiated.² The subsequent “3” is then retrieved from memory. At that moment the control production that protects memory consolidation is triggered, blocking target detection until memory consolidation is finished. At the end of consolidation, another control production restarts target detection (“done consolidation”), after which target detection resumes. However, in the intermediate period several visual inputs have been ignored, most notably the “B” target.

Fig. 2b shows an example of Lag-1 sparing. In that case, the production that recognizes the second target supersedes the control production that protects consolidation. As a consequence, the target detection skill holds on to T2 (because of politeness) during T1 consolidation. This means that in practice only distractors trigger an AB. There is variability due to noise in many aspects of the model: the time the visual module needs to identify an item, the time declarative memory needs to retrieve it and the time the imaginal system needs to consolidate a target. This explains why the model does not always behave in the same way on a given lag. For example, Fig. 2c shows a trace of a Lag-3 trial in

² In this case, the two rules from Fig. 1, “A is a target” and “Store target A”, have been collapsed into a single production rule. This is achieved by the production compilation learning mechanism in ACT-R (Taatgen & Anderson, 2002), which we will not detail here because it only plays a minor role in the model.

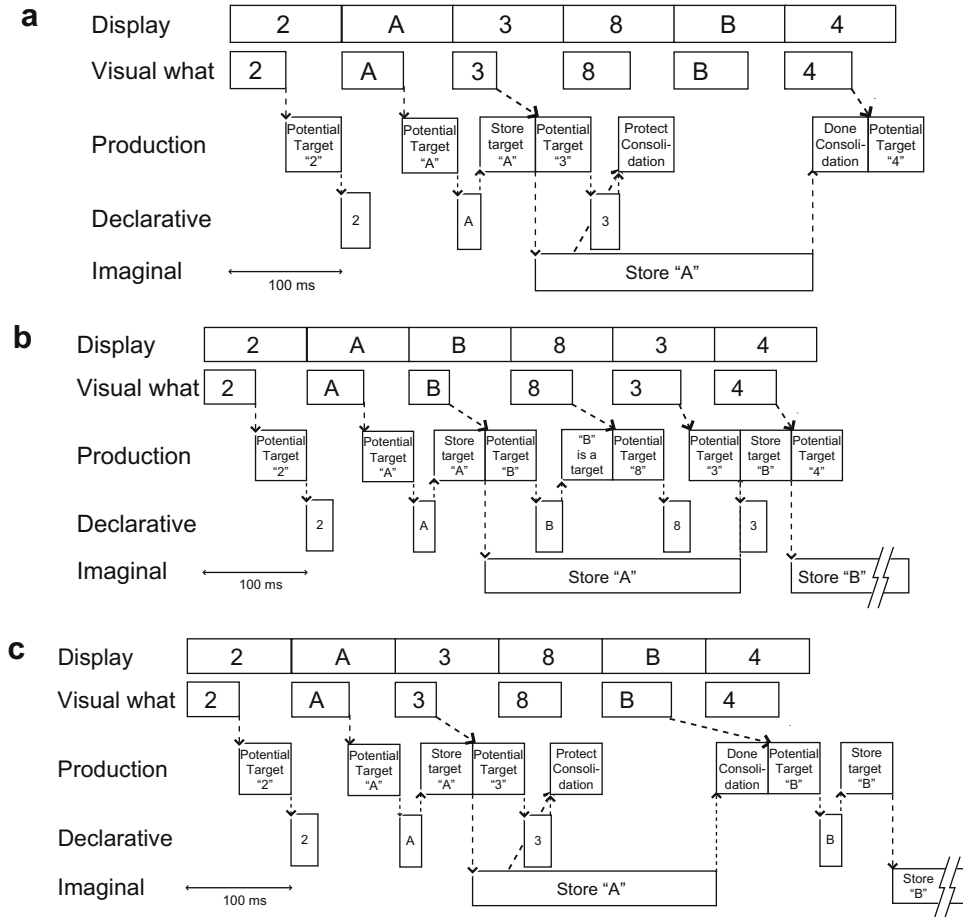


Fig. 2. Example of a model trace for (a) a Lag-3 blink, and (b) Lag-1 sparing (c) Lag-3 without blink.

which T2 is correctly identified. The difference with the trace in Fig. 2a is that in Fig. 2c the "B" stimulus is still in the visual buffer when the "done consolidation" production finishes, while in Fig. 2a it is already replaced by the subsequent "4".

Details of the model can be found in Appendix A, which lists the production rules in the model, and the parameter values that are used in the model throughout the article. The function of parameters is to set processing durations in each of the modules. We will show how the model can explain several existing phenomena in the literature, and then proceed with the model's new predictions.

2.1. Model of the standard blink

Fig. 3 shows the fit of the model to results of an experiment reported by Taatgen, Juvina, Herd, Jilk, and Martens (2007). In this experiment, 37 participants performed the RSVP task in which letters had to be detected among digit distractors. The Lag in the two-target trials varied from 1 to 9. Eight participants were classified as non-blinkers based on their T2|T1 accuracy, and 29 as blinkers. The graph shows the data of the blinkers: we will discuss the non-blinkers in a later section. The model produces a good fit of the data, with one exception: its T1 accuracy on Lag-1 is identical to the other Lags, while in the data performance is worse. This is related to several shortcomings of the model regarding Lag-1 performance.

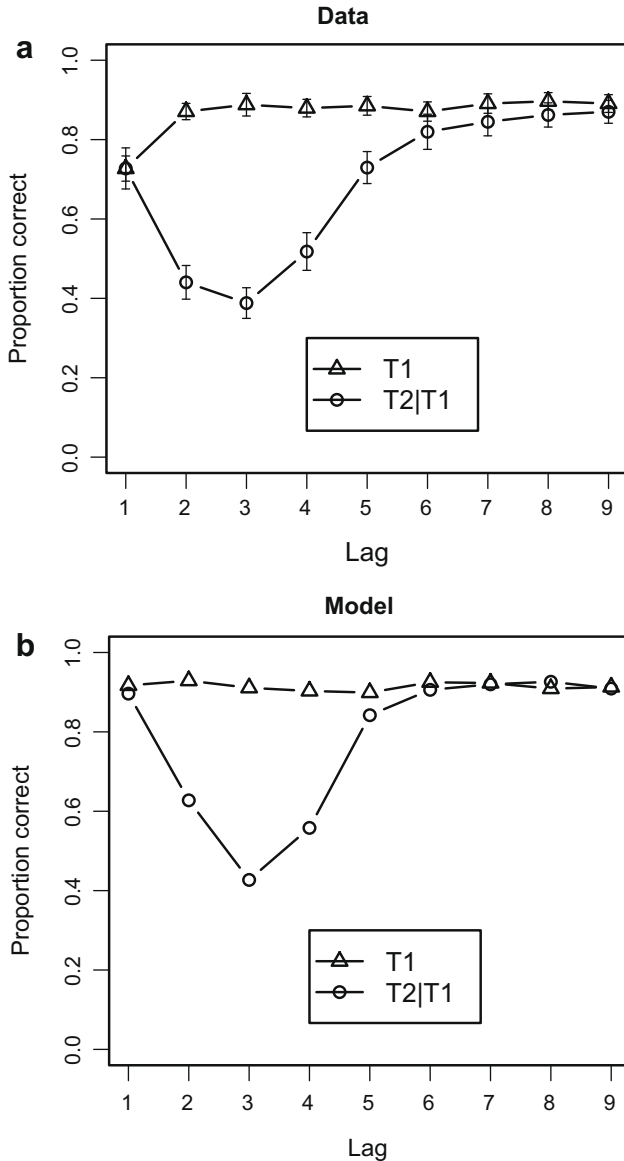


Fig. 3. T1 accuracy and T2 accuracy given T1 correct for the standard blink task, data and model.

In order to remedy these problems, we have developed a neural network version of the visual module that adequately captures several Lag-1 issues. Because modeling the various Lag-1 issues is relatively unrelated to the main issue we address here, we will only briefly discuss that version of the model here.

2.2. Neural network visual module

A property of Lag-1 performance is that the two targets are often reported in reverse order. The current model cannot explain this, because the default ACT-R visual module identifies single discrete items. Instead, participants often report that they see targets in Lag-1 trials as superimposed, and

can therefore not determine the order. To model this aspect of the task, we replaced the standard ACT-R visual module with a Leabra neural network (O'Reilly & Munakata, 2000). The details are explained in Taatgen et al. (2007). In the neural visual module, each potential item (letter or digit) is represented by an output cell. Because of the speed of the input representation, two output cells can be active at the same time: one for the present stimulus, and the other for the previous stimulus, which activation still lingers in the network. In the case of two active targets, the production rule that samples the visual input cannot determine the order, and randomly guesses which was first. This version of the model correctly fits the reversed report of targets on Lag-1 trials and also cases in which T1 is not reported, but T2 is reported correctly. It can also consolidate T1 and T2 at the same time on Lag-1 trials, which is more consistent with ERP data of Lag-1 (Craston, Wyble, Chennu, & Bowman, *in press*).

2.3. Model of longer sequences of targets

Di Lollo et al.'s (2005) experiment, which we discussed in the introduction, contrasted regular Lag-2 trials (T1–Distractor–T2, the varied condition) with trials in which three subsequent targets (without intervening distractor) had to be reported (T1–T2–T3, the uniform condition). Fig. 4 shows the results of that experiment: whereas the varied condition shows the regular blink effect, the uniform condition does not. Fig. 4 also shows the results of our threaded cognition model: although the model slightly overpredicts accuracy in the uniform condition, it nevertheless reproduces the basic finding. The reason why the model does not exhibit the blink is that in a sequential stream of targets, the control production that suspends target detection never fires: the production that detects a target has a higher utility value, and therefore always takes precedence over it.

2.4. Regular versus speeded presentation

Bowman and Wyble (2007) contrasted a regular RSVP experiment (with 100 ms/item) with a version in which the presentation rate was increased to 50 ms/item (also see Martens et al., 2006). The results (Fig. 5a) are plotted against time instead of Lag. This means that the 100 ms target onset asynchrony corresponds to Lag-1 in the 100 ms/item condition, but to Lag-2 in the 50 ms/item condition. Even though the overall accuracies in the 50 ms/item condition are much lower, the blink occurs in the same time interval of 200–500 ms after T1. In the model, the increased presentation rate produces a lower overall accuracy, because the visual module has trouble processing input at this rate. However, the control production and the time the imaginal module needs to consolidate targets are independent of the presentation rate, allowing it to fit the data correctly (Fig. 5b).

2.5. Insertion of blanks in the RSVP stream

Both Bowman and Wyble (2007) and Shih (2008) list other findings and effects associated with the AB that models should explain and replicate. Some of these have already been discussed in this paper: the basic blink and Lag-1 sparing are part of the basic model fit. The neural network model addresses specific issues with Lag-1: reduced accuracy of T1, swaps between T1 and T2, and cases in which only T2 is reported. Another known effect is that if there is a blank in the T1+1 position instead of a distractor, the AB is decreased, but not when it is in the T1+2 position (Chun & Potter, 1995; Raymond et al., 1992). This effect can be readily accounted for in our model, because the control production that suspends target detection is partially triggered by a declarative retrieval. A blank does not initiate any retrieval, resulting in a postponement of the AB.

Fig. 6 shows the data from Chun and Potter (1995), and the results of the model. The basic finding, T1+1 Blank decreases the blink while T1+2 Blank does not, is reproduced. A related effect is that if T2 appears at the end of the stream (T2 unmasked), there is no blink (Giesbrecht & Di Lollo, 1998). For the model this is a trivial case, because the last perceived item remains in the visual buffer, giving the model ample time to process it.

Lastly, several studies have shown priming effects in the RSVP task (e.g., Chua, Goh, & Hon, 2001; Shapiro, Driver, Ward, & Sorensen, 1997). Although we have not explored effects of priming in this article, the threaded cognition model should be able to deal with priming effects. All items in the vi-

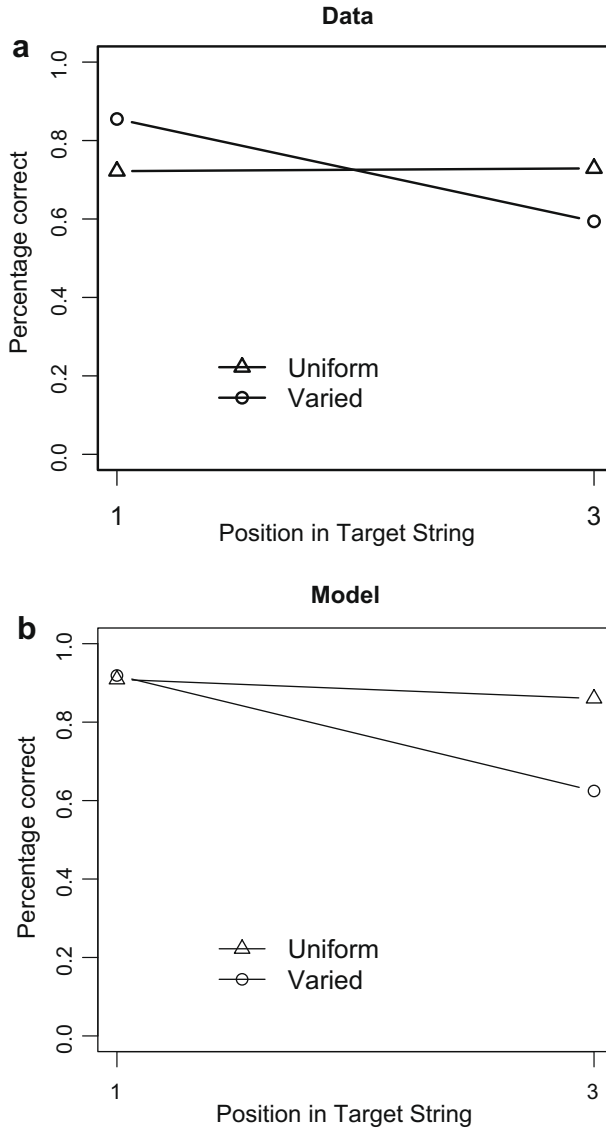


Fig. 4. Comparison between data from Di Lollo et al. (2005) and the model prediction. Data figure adapted from “The attentional blink: Resource depletion or temporary loss of control?” by V. Di Lollo et al., 2005, *Psychological Research*, 69, p. 194. Copyright 2005 Springer Verlag. Adapted with permission.

sual stream are attended by the visual module, and therefore impact the activation levels of the corresponding memory traces in declarative memory (this is an automatic feature of ACT-R). A possible effect of priming in the model is a faster declarative retrieval time, leading to more accurate recognition of primed items.

3. Non-blinkers and decreased control

Although the AB phenomenon is very robust, a small proportion of participants, referred to as non-blinkers, does not show any attentional blink at all (Feinstein et al., 2004; Martens & Valchev, 2009;

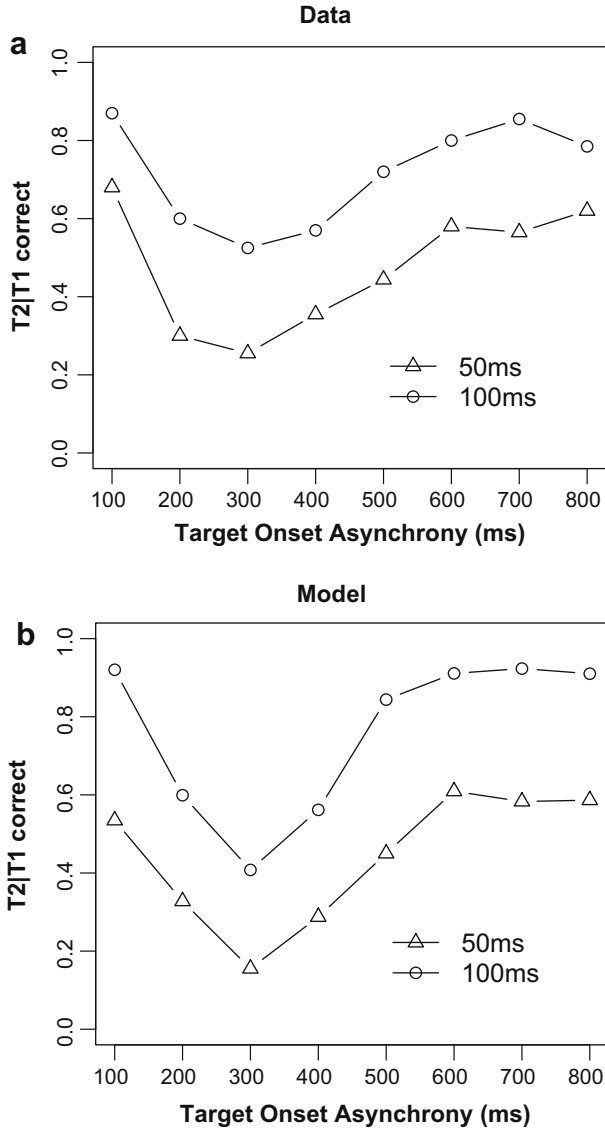


Fig. 5. Comparison between data from [Bowman and Wyble \(2007\)](#) and the model prediction. Data figure adapted from “The simultaneous type, serial token model of temporal attention and working memory”, by H. Bowman and B. Wyble, 2007, *Psychological Review*, 114, p. 55. Copyright 2007 by the American Psychological Association. Adapted with permission.

[Martens et al., 2006, 2009](#)). The interesting question is what is different about these individuals. An ERP study by [Martens et al. \(2006\)](#) shows that detection of a target correlates with a P300 peak in the parietal area, because it is present when a target is detected and absent when it is not. They therefore associate the P300 with memory consolidation. The study also shows that non-blinkers' P300s occur earlier than those of blinkers, especially those triggered by the second target.

Non-blinker behavior can be explained by leaving out the control production that blocks target detection during memory consolidation. Analogous to Lag-1 sparing this does not lead to a disruption of T1 consolidation, because the politeness policy of threaded cognition postpones consolidation of T2

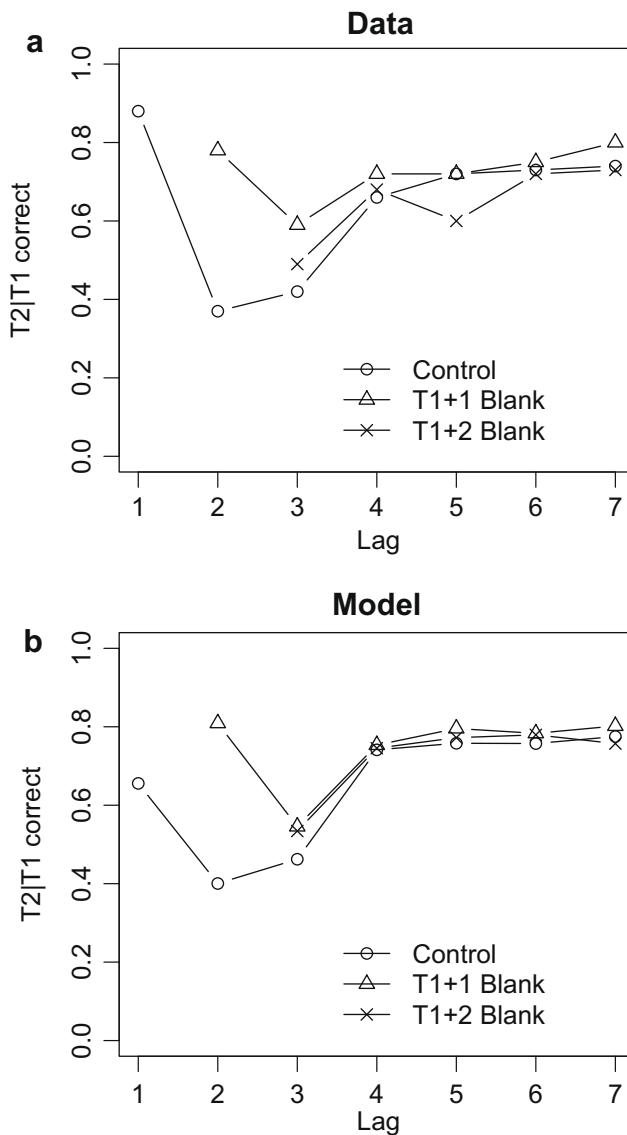


Fig. 6. Comparison between data from Chun and Potter (1995) and the model prediction. Data figure adapted from “A 2-Stage Model for Multiple-Target Detection in Rapid Serial Visual Presentation”, by M.M. Chun and M.C. Potter, 1995, *Journal of Experimental Psychology-Human Perception and Performance*, 21, p. 114. Adapted with permission.

until T1 has been consolidated. Fig. 7 shows a trace of how the absence of this production leads to the absence of an AB. After the memory-consolidation process has started, the target detection process can continue without any interference. The trace shows how target detection and memory consolidation can operate in almost perfect parallelism for as long as control processes do not interfere. It also demonstrates why memory consolidation for the second target starts earlier for non-blinkers than for blinkers. With blinkers, target detection is halted during memory consolidation, and determining whether a second potential target is a target only resumes after the first target has been consolidated (Fig. 2c). Non-blinkers, on the other hand, keep processing potential targets during memory consolidation, and can therefore immediately start consolidating a second target once the first target is done.

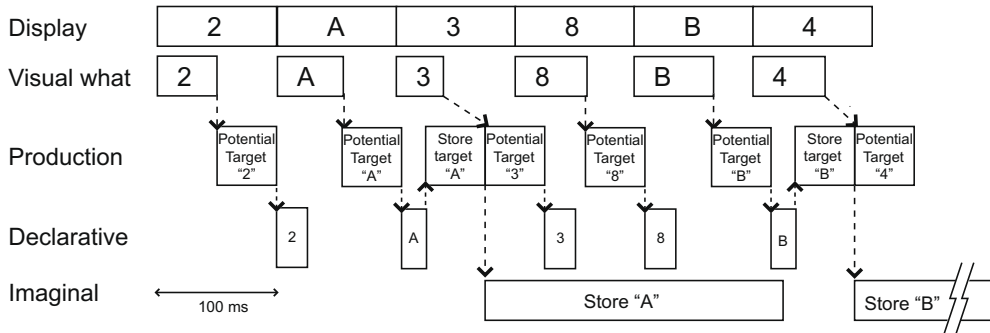


Fig. 7. Model trace of a non-blinker on Lag-3.

This is consistent with the ERP results from Martens et al. (2006) mentioned earlier. Moreover, the parietal activation found by Martens et al. fits ACT-R's theory that imaginal-buffer activity is localized in the parietal cortex (see Anderson, 2007, for an elaborate discussion the localization of ACT-R's modules in the brain). Fig. 8 shows a comparison between the eight participants from the Taatgen et al. (2007) study that we classified as non-blinkers and the model fit.

The absence of the control production that protects consolidation can explain why some individuals do not exhibit an AB in RSVP tasks. We do not have a good explanation why non-blinkers have a different control strategy, but we can try to induce a control scheme similar to that of the non-blinkers by using a manipulation that decreases control. Several researchers have found that a reduction in the level of attention on the RSVP task attenuates the AB. Experimental manipulations can decrease the amount of blink people exhibit, for example if the stimuli are presented in a star field (Arend et al., 2006), when music is played in the background (Olivers & Nieuwenhuis, 2005), or when participants receive instructions to focus less on the task (Olivers & Nieuwenhuis, 2006). Although these manipulations have been reasonably successful, they have not been altogether consistent: Olivers and Nieuwenhuis (2006, footnote 1) mention that the addition of music has not produced the effect consistently, and our own attempts of replicating the effect of a star field in the background (Arend et al., 2006) have not been successful.

The model can explain a decrease in attentional blink due to distraction by assuming that the procedural memory resource becomes overloaded. According to threaded cognition, multiple tasks can be done in parallel as long as the resources they need do not overlap in time. However, adding additional tasks will put additional demands on the procedural resource (and possibly other resources as well). In normal situations, overloading a resource leads to slower performance, because the different tasks or threads have to wait until it is their turn to use the resource. However, in an RSVP task there is no time to wait, so overloading the procedural resource will lead to dropping production activations. The productions that are dropped first are the productions with the lowest utility, which in our case is the control production that leads to the blink. Both the star field and music manipulations are tasks that require production activations, unless they are successfully ignored. This may explain the fragility of some of the results involving distraction: if a distraction does not provide enough production activations to supersede the control production, no decrease in the blink will follow.

With this in mind, we designed a new manipulation in the form of a secondary task, which forces participants to divert attention to it. The added advantage is that it is much easier to model the demands of this task than the demands produced by listening to music or looking at a star field.

4. Experiment 1

Participants performed two tasks, an AB task during which 0–2 letters had to be identified within an RSVP stream of digits, and a modified version in which a concurrent red³ dot detection task had to

³ For interpretation of color in Fig. 10, the reader is referred to the web version of this article.

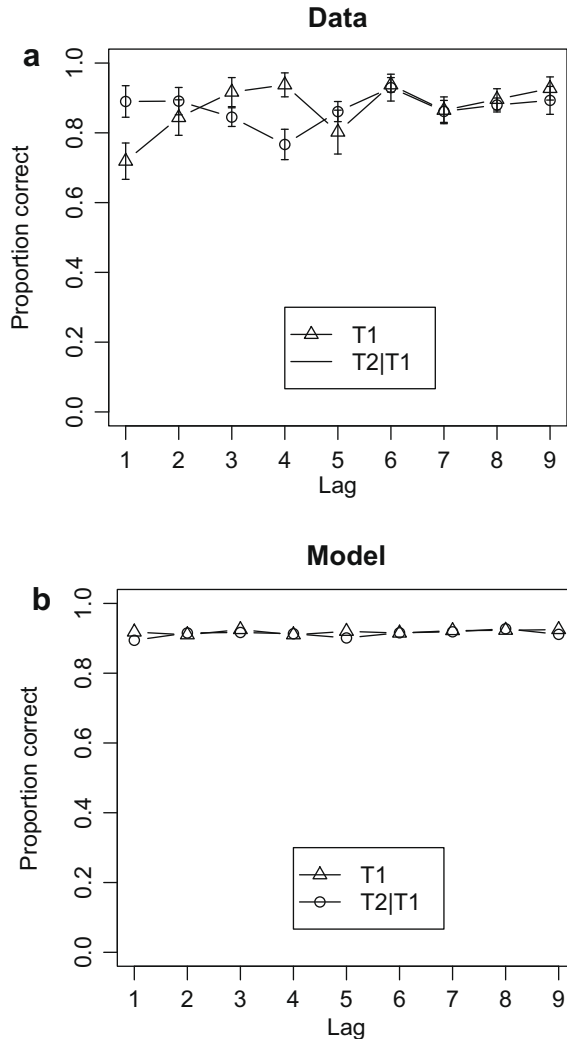


Fig. 8. T1 accuracy and T2 accuracy given T1 correct for the non-blinkers, data and model.

be carried out along with the AB task. In the red dot detection task a moving grey dot is presented in peripheral vision concurrently with the central RSVP stimuli. Participants have to detect whether or not the dot briefly turns red at some time during the trial.

4.1. Method

4.1.1. Participants

Thirty volunteers from the University of Groningen community (aged 18–27, mean = 21.2, with normal or corrected-to-normal visual acuity) participated and received payment of € 12. Half of the participants (standard First group) performed the AB task without red dot detection first, followed by the AB task with red dot detection in a second session (at least a week later). The other half (Dot First group) performed the AB task with dot detection first, followed by the AB task without dot detection in a second session. Participants were randomly assigned to each group.

4.1.2. Stimuli and apparatus

Stimuli were black (17.9 cd/m^2) digits (excluding 1 and 0) and uppercase consonants (excluding 'Q', 'V', and 'Y') presented on a white background (88.0 cd/m^2) in a bold 12-point Courier New font subtending 0.3 by 0.4° of visual angle at a viewing distance of approximately 50 cm on a 17-in. monitor. In the AB task with red dot detection, grey (40.2 cd/m^2) and red (28.8 cd/m^2) dots with a diameter of 10 pixels were used. The generation of stimuli and the collection of responses were controlled using E-prime 1.1 software (Schneider, Eschman, & Zuccolotto, 2002) running under Windows XP on a PC with a 2.8 GHz processor.

4.1.3. Procedure

Each task consisted of a practice block of 24 trials and three testing blocks of 96 trials each. At the start of each block, four additional warm-up trials were provided that were excluded from the analyses. A short break was given after each block.

On each trial, a fixation cross was presented in the middle of the screen. At the bottom of the screen, participants were prompted to press the space bar to initiate the trial. In the AB task without red dot detection, the message disappeared when the space bar was pressed, and 100 ms later the RSVP stream was presented. For the AB task with red dot detection, a grey dot was sequentially presented together with each RSVP item for 90 ms, randomly starting at one of 39 possible peripheral positions, moving in clockwise rotation (skipping two positions with each presentation) in a radius of 11.3° from the middle of the screen. In 25% of the trials one of these 16 grey dots was randomly replaced by a red dot.

In two thirds of the trials of both AB tasks, two targets were embedded in the stream (dual-target trials), in one sixth of the trials only one target letter was present (single-target trials), and in one sixth of the trials no targets were present (no-target trials). In dual- and single-target trials, T1 was always presented as the sixth item in the stream. In dual-target trials, T2 was the first, second, third, or eighth item following T1 (i.e., it was presented at Lags 1, 2, 3, or 8, respectively). These specific Lags were chosen on the basis of the literature and previous work in our laboratory. T2 is likely to be "blinked" (i.e., not identified) at Lags 2 and 3, whereas at Lags 1 and 8 little or no reduction in T2 accuracy is usually observed. Target letters were randomly selected with the constraint that T1 and T2 were always different letters. Digit distractors were randomly selected with the constraint that no single digit was presented twice in succession.

After the RSVP stream was presented, the screen was cleared and participants were prompted by a message at the bottom of the screen to type the letters they had seen using the corresponding keys on the computer keyboard. Participants were instructed to take sufficient time in making their responses to ensure that typing errors were not made. If a letter was not seen, the space bar was to be pressed instead. Participants were encouraged to type their responses in the order in which the letters had been presented, but responses were accepted and counted correct in either order. No feedback on AB performance was given. In the AB task with red dot detection, participants were subsequently prompted to indicate whether a red dot had been presented or not by pressing 'j' on the keyboard for yes or 'n' for no. Feedback on red dot detection performance was provided. Participants initiated the next trial by pressing the space bar. The AB task without red dot detection (standard task) was completed in approximately 35 min, and the AB task with red dot detection (dot task) was completed in approximately 40 min.

4.2. Results

For the dot task trials, only the trials were included in the analysis in which the dot did not turn red (75% of the trials). In the dot task, accuracy for detecting the dot was 59.2% for the standard task first group, and 47.6% for the dot task first group, which was not different ($p = 0.28$). Further analyses were restricted to dual-target trials.

Table 1 shows the accuracies for both T1, and T2 given that T1 was correct for each condition and block. An analysis of variance (ANOVA) of these accuracies with condition and block as factors only reveals a main effect of block, indicating that participants get better with practice (for T1:

$F(1,28) = 5.80$, $MSE = 0.015$, $p = 0.023$, for $T2|T1$: $F(1,28) = 7.77$, $MSE = 0.032$, $p = 0.009$). This means that the addition of the dot task does not have an impact on overall accuracy of the RSVP task.

Fig. 9 shows the accuracies for the different lags. These graphs suggest that the dot manipulation only has an impact on the T2 accuracy in the standard first condition. In order to assess this, we calculated blink magnitude, similar to for instance Colzato, Spapé, Pannebakker, and Hommel (2007), by taking the difference between average T2|T1 accuracy on Lags 1 and 8, and Lags 2 and 3, respectively. The bottom two rows in Table 1 show the magnitudes for the two conditions and blocks. An analysis of

Table 1
Mean accuracies (and standard deviations) of T1, T2|T1, and blink magnitude by block and condition in Experiment 1.

		Block 1	Block 2
T1	Standard first	0.83 (0.12)	0.85 (0.11)
	Dot first	0.83 (0.11)	0.87 (0.11)
T2 T1	Standard first	0.65 (0.22)	0.72 (0.20)
	Dot first	0.72 (0.18)	0.74 (0.20)
Blink	Standard first	0.34 (0.18)	0.09 (0.14)
Magnitude	Dot first	0.12 (0.13)	0.13 (0.14)

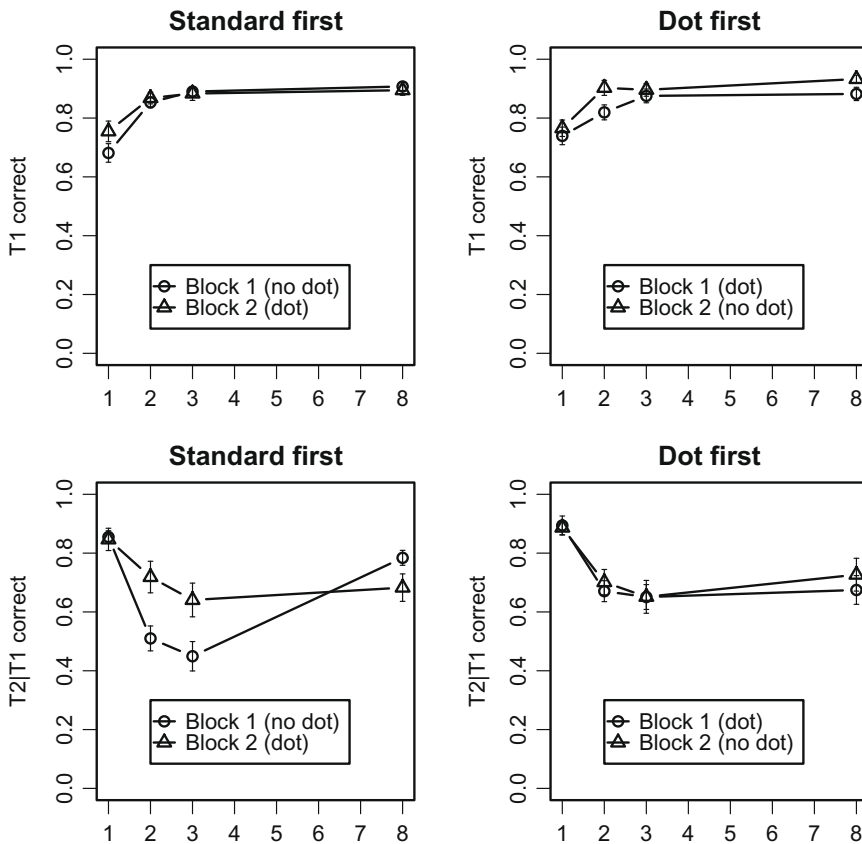


Fig. 9. Mean percentage correct report of T1 and T2|T1 in Experiment 1 for the Standard First and Dot First groups as a function of lag and session.

variance of the blink magnitude with block and condition as factors reveal a main effect of block, $F(1,28) = 19.3$, $MSE = 0.23$, $p < 0.001$, an interaction between block and condition, $F(1,28) = 21.4$, $MSE = 0.26$, $p < 0.001$, and a weak main effect of condition, $F(1,28) = 3.39$, $MSE = 0.109$, $p = 0.08$. Apparently the dot manipulation has an impact on the amount of blink, but in addition there is an overall effect of block. There are two possible explanations for these results. The first is that the main effect of block is due to learning, which causes the effect of the dot and of learning to cancel each other in the second block of the dot-first condition. The second explanation is that the decreased control induced by the dot manipulation in the dot-first conditions persists even when the dot task is taken away in the second block. One of the goals of Experiment 2 will be to see which explanation is correct.

4.3. Model

The addition of a secondary task is relatively easy to model within the framework of threaded cognition: a separate model of dot detection is constructed, which is then combined with the existing AB model. The model for dot detection is very simple: each time the dot moves, a production rule activates that checks its color (Fig. 10). The color of the dot can be checked in peripheral vision, and can therefore be accessed independently of the visual stream that identifies letters. Whenever a red dot is detected, a memory consolidation step as in Fig. 1c is initiated to store the fact that the dot has been seen.

With the addition of the dot task, the production resource becomes overloaded and starts acting as a bottleneck, while it did not earlier on. Fig. 11 shows an example for Lag-3. In this example, the dot stays grey, but requires production rules to monitor it each time it moves. Consequently, the production rule that protects the consolidation process by suppressing target detection never fires. The model therefore behaves more like a non-blinker. We will discuss the model fits in the context of Experiment 2.




Fig. 10. Module activation produced by the dot-detection task.

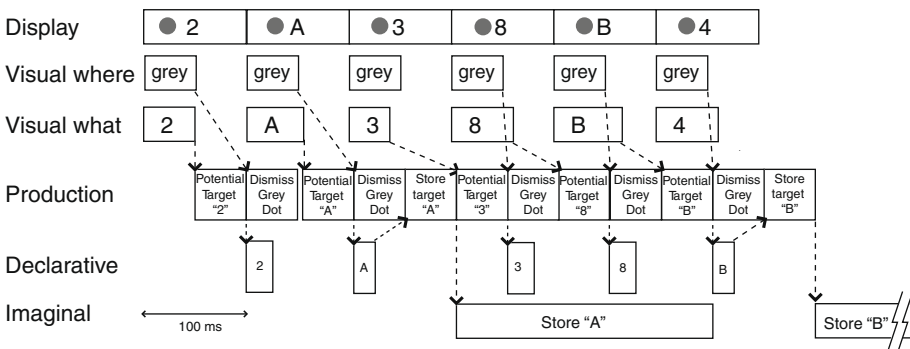


Fig. 11. Example model trace of Lag-3 sparing due to the addition of the dot task.

4.4. Discussion

Although the results suggest that the addition of the secondary dot task indeed attenuates the AB, there are still two possible interpretations. One interpretation is that the dot task decreases the magnitude of the AB, and that this effect lasts even when the dot task is removed. The second explanation is that the results are produced by a combination of the dot task and a learning effect: when participants start with the standard condition, and then move to the dot condition, their performance improves both due to the effect of the dot task and the effect of learning. When they start in the dot condition and then move to the standard condition, the effects of learning and moving to the “harder” no-dot condition cancel each other out. The model is consistent with the second explanation: as soon as the dot task is taken away, the production that protects memory consolidation can once more fire, reintroducing the blink.

5. Experiment 2

Experiment 2 was designed to control for the effects of learning in the task in order to decide which of the two explanations from Experiment 1 is correct, and to replicate the finding that the dot task indeed decreases the blink (given that some of the other manipulations have not always been successful in replicating blink-attenuating effects). A second goal is to show that the model can also predict increases in the amount of blink. The decrease in blink caused by the dot task, according to the model, is due to overloading the procedural resource, which makes it less likely that the production rule that causes the blink will activate. A similar result can be expected if other tasks that require production rules are added. Increasing the load on one of the other resources can have different effects. For example, increasing the load on the imaginal resource will increase the blink, because the imaginal resource will be unavailable more often, giving the control production more opportunities to activate and produce a blink. In order to increase imaginal load, we designed a new version of the task in which the first target is rotated 180°. Recognizing that a rotated T1 is a target is not slower, and may even be faster because the fact that it is rotated identifies it as a target (Heil, 2002). However, identifying objects as opposed to determining their categories does slow down with rotation (Dickerson & Humphreys, 1999), especially when objects are not symmetric (Leek & Johnstons, 2006). Whether or not this process is really mental rotation, or some other form of mapping, it is generally associated with the parietal cortex (Alivisatos & Petrides, 1997; Heil, 2002), the area that corresponds to the imaginal resource (Anderson, 2007). According to the model, adding rotation to the task will increase the load on that resource, and therefore increase the AB.

Experiment 2 consists of four blocks, in which we add the experimental manipulation (control, dot or T1 rotation) in block 3. All the other blocks use the standard blink experiment, allowing a measure of the learning between blocks 1 and 2, a measure of the effects of the manipulation in block 3, and a measure to what extent effects persist after the manipulation is taken away by comparing blocks 3 and 4, and blocks 2 and 4.

5.1. Method

5.1.1. Participants

Sixty-three volunteers from the Carnegie Mellon University community (aged 18–47, mean = 22, with normal or corrected-to-normal visual acuity) participated and received a payment of \$10. Participants were randomly assigned to one of three groups: Control, Dot, or Rotated. Each of the groups performed four blocks, where all blocks were standard AB tasks, except for the third block, where the Control group received the standard AB task, the dot group the AB task with dot detection, and the Rotated group the AB task with a rotated first target.

5.1.2. Stimuli and apparatus

Stimuli were black digits (excluding 1 and 0) and letters from the following set: A, E, C, D, G, J, K, L, M, N, R, T, V, W. They were presented on a grey background in 30-point Lucida Grande. In the AB task

with red dot detection, grey and red dots with a diameter of 10 pixels were used. The generation of stimuli and the collection of responses were controlled by a specific application written in the XCode programming environment on an Apple Macintosh computer.

5.1.3. Procedure

Each block consisted of 112 trials, 16 of which had no targets, 16 had one target, and 80 had two targets. Before block 1 and block 3 a practice block of 6 trials was given, with 2 no-target, 2 one-target and 2 two-target trials. On each trial, a fixation dot was presented in the middle of the screen for 500 ms, after which the RSVP stream was presented. Each stream consisted of 15 items that were presented at a rate of 100 ms/item. In the single- and dual-target trials, the first target was presented in position 5, and in the dual-target trials the second target at position 6, 7, 8, 11 or 13, producing trials with Lags 1, 2, 3, 6 and 8, respectively. In the block with rotated T1 trials, T1 was rotated 180°. In the block with red dot detection, a grey dot was presented together with each RSVP item for 100 ms. As in Experiment 1, in 25% of the trials one of the 15 grey dots was randomly replaced by a red dot. After presentation the participants were prompted to enter the targets they had seen, and, in the dot condition, were asked to indicate whether they had seen a red dot by pressing 'y' or 'n'. No feedback was given on either aspect of the task. The program proceeded to the next trial after the participant had entered the appropriate responses. The whole experiment was completed in approximately 50 min.

5.2. Results

Table 2 shows the average accuracies for T1 and for T2|T1 taken over all lags. The accuracies in the dot condition have been split into trials in which the dot turned red (25% of the trials), and trials in which the dot stayed gray (75% of the trials). The table shows that the addition of a dot has little impact on the accuracy as long as it does not turn red. If the dot turns red, on the other hand, accuracies are significantly lower: $t(20) = 2.3$, $p = 0.033$ for T1 accuracy, and $t(20) = 3.95$, $p < 0.001$ for T2|T1 accuracy. In the dot group, accuracy for detecting the dot in trials where the dot turned red was 93%, which is considerably higher than in Experiment 1 (59.2%). In the following analyses, we will only include the trials in the dot condition in which the dot stayed gray.

An analysis of variance of the average T1 accuracy with condition and block as factors reveals an effect of block, $F(3, 160) = 2.94$, $MSE = 0.018$, $p = 0.035$, and an interaction between block and condition, $F(6, 160) = 2.62$, $MSE = 0.017$, $p = 0.019$, but no main effect of condition, $F(2, 60) = 2.24$, $MSE = 0.064$, $p = 0.12$ (see also Fig. 13). All these effects are due to the rotated condition, because the same analysis with just the control and dot condition yields no effects (all $F_s < 1$). Both the main effect of block and the interaction are stronger in a comparison between the control and the rotated condition, $F(3, 120) = 4.58$, $MSE = 0.026$, $p = 0.005$ and $F(6, 120) = 4.64$, $MSE = 0.026$, $p = 0.004$, respectively. This can probably be explained by the fact that a rotated letter is slightly harder to identify, reducing accuracy. A similar analysis of variance of the T2|T1 accuracies shows a main effect of block, $F(3, 180) = 6.21$, $MSE = 0.047$, $p < 0.001$, and an interaction between block and condition, $F(6, 180) = 4.28$, $MSE = 0.032$, $p < 0.001$, but no main effect of condition, $F(2, 60) = 1.11$, $MSE = 0.087$, $p = 0.34$. The interaction is again due to poorer performance in the rotated condition: a comparison

Table 2

Mean accuracies (and standard deviations) of T1 and T2|T1 by block and condition in Experiment 2.

		Block 1	Block 2	Block 3	Block 4
T1	Control	0.86 (0.086)	0.87 (0.081)	0.87 (0.080)	0.87 (0.063)
	Dot (no red dot)	0.85 (0.085)	0.89 (0.089)	0.88 (0.083)	0.86 (0.12)
	Dot (red dot)			0.79 (0.12)	
	Rotate	0.84 (0.085)	0.85 (0.11)	0.84 (0.11)	0.80 (0.12)
T2 T1	Control	0.64 (0.17)	0.68 (0.17)	0.72 (0.17)	0.71 (0.10)
	Dot (no red dot)	0.66 (0.13)	0.73 (0.15)	0.72 (0.22)	0.74 (0.16)
	Dot (red dot)			0.61 (0.29)	
	Rotate	0.65 (0.12)	0.70 (0.10)	0.58 (0.18)	0.67 (0.16)




Fig. 12. Blink magnitude for the four blocks and three conditions in Experiment 2. The experimental manipulation (dot or rotated), was only administered in block 3.

between the rotated and the control condition has a significant interaction, $F(3,120) = 7.75$, $MSE = 0.056$, $p < 0.001$, while a comparison between the dot and the control condition shows no effect, $F(3,120) < 1$. We can therefore conclude that the dot manipulation has no effect on the overall accuracies as long as the dot stays gray, and that the rotated condition decreases accuracy on both T1 (which is rotated) and T2 (which is not rotated).

5.2.1. Impact on the blink magnitude

To analyze the impact of the different manipulations on the amount of blink, we have calculated, analogous to Experiment 1, the difference between the average T2|T1 on trials for Lags 1, 6, and 8, and the average T2|T1 for Lags 2 and 3. Fig. 12 plots these difference scores, which gives an indication of the size of the blink, for the four blocks and three conditions. The Figure suggests an overall learning effect, and an impact of condition on block 3. Fig. 13 details the impact of condition in block 3, and shows that the dot manipulation decreases the blink, and the rotated manipulation increases it. An analysis of variance with block and condition as factors confirms both a learning effect and an effect of condition: there is a main effect of block, $F(3,180) = 18.2$, $MSE = 0.27$, $p < 0.001$ and an interaction between block and condition, $F(6,180) = 5.6$, $MSE = 0.084$, $p < 0.001$, but no main effect of condition $F(2,60) = 1.25$, $MSE = 0.27$, $p = 0.29$. Pair wise comparison between the dot and the control condition also shows a significant interaction between condition and block, $F(3,120) = 4.72$, $MSE = 0.067$, $p = 0.004$. The interaction is, however, not significant between the control condition and the rotated condition, $F(3,120) = 1.50$, $MSE = 0.023$, $p = 0.22$. We can therefore conclude that the dot manipulation has a significant effect on the amount of blink when compared to the control condition, but that the difference with the rotated condition does not reach significance.

5.2.2. Within-group effects

Apart from between-group comparisons, we can also compare the magnitude of the blink within each group. Within-subject t -tests showed that in the dot condition there was a marginal decrease in difference score between blocks 1 and 2, probably due to learning, $t(20) = -1.79$, $p = 0.09$, a significant decrease between blocks 2 and 3, due to the introduction of the dot task, $t(20) = -4.84$, $p < 0.001$, and a significant increase between blocks 3 and 4, due to the removal of the dot task, $t(20) = 3.31$, $p = 0.004$. However, there was no difference between blocks 2 and 4 ($t(20) = -0.83$, $p = 0.42$), indicating that any improvement between blocks 2 and 3 is not due to learning.

In the rotated condition there was a decrease of blink magnitude between blocks 1 and 2, $t(20) = -3.14$, $p = 0.005$, an increase between blocks 2 and 3, $t(20) = 2.46$, $p = 0.023$, and a decrease be-

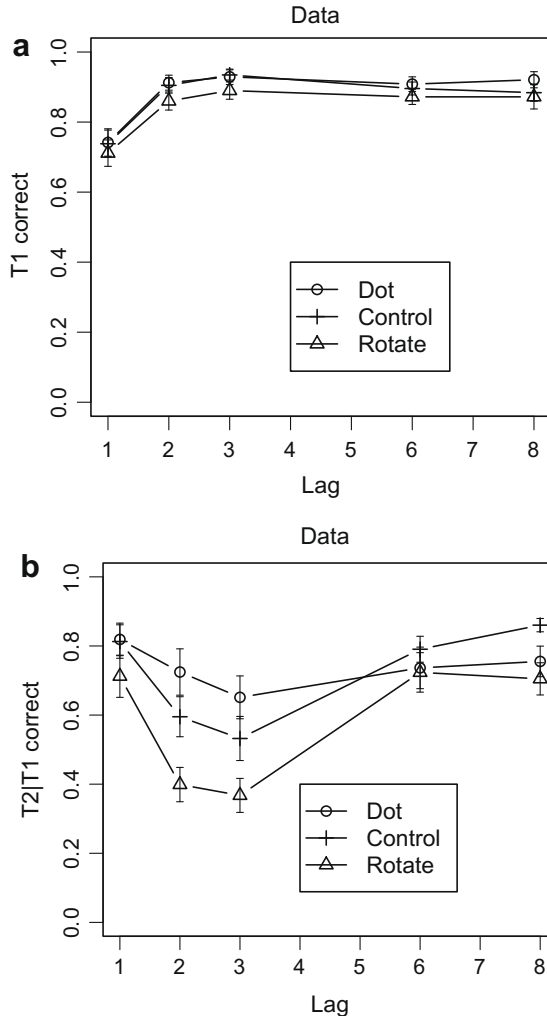


Fig. 13. Mean proportion of correct T1 and correct T2 given T1 correct in block 3 of Experiment 2.

tween blocks 3 and 4, $t(20) = -3.28$, $p = 0.004$. Again, there was no difference between blocks 2 and 4, $t(20) = -1.51$, $p = 0.15$. From this we can conclude that there is a significant impact of the rotated letter on the blink magnitude, even though the contrast with the control group did not yield a significant result. Finally, in the Control group there was only a significant decrease in blink between blocks 1 and 2, $t(20) = -2.85$, $p = 0.010$, but not between blocks 2 and 3, and blocks 3 and 4, $t(20) = -0.70$, $p = 0.49$ and $t(20) = -1.91$, $p = 0.07$, respectively.

5.2.3. Analysis of the trials with a red dot

Fig. 14 shows further details on block 3 of the dot condition, in which trials in which the dot turned red are compared to trials in which the dot stayed gray. There is an overall decrement in accuracy on both T1 and T2|T1 in trials in which the dot turns red (as was already noted at the start of the results section). In the case of T1 there is no relationship to the lag. In the case of T2|T1, there is a trend towards a larger difference impact of the red dot on the longer lags ($F(4, 171) = 2.29$, $MSE = 0.80$, $p = 0.06$). A possible explanation for this is that the dot itself can sometimes produce an attentional

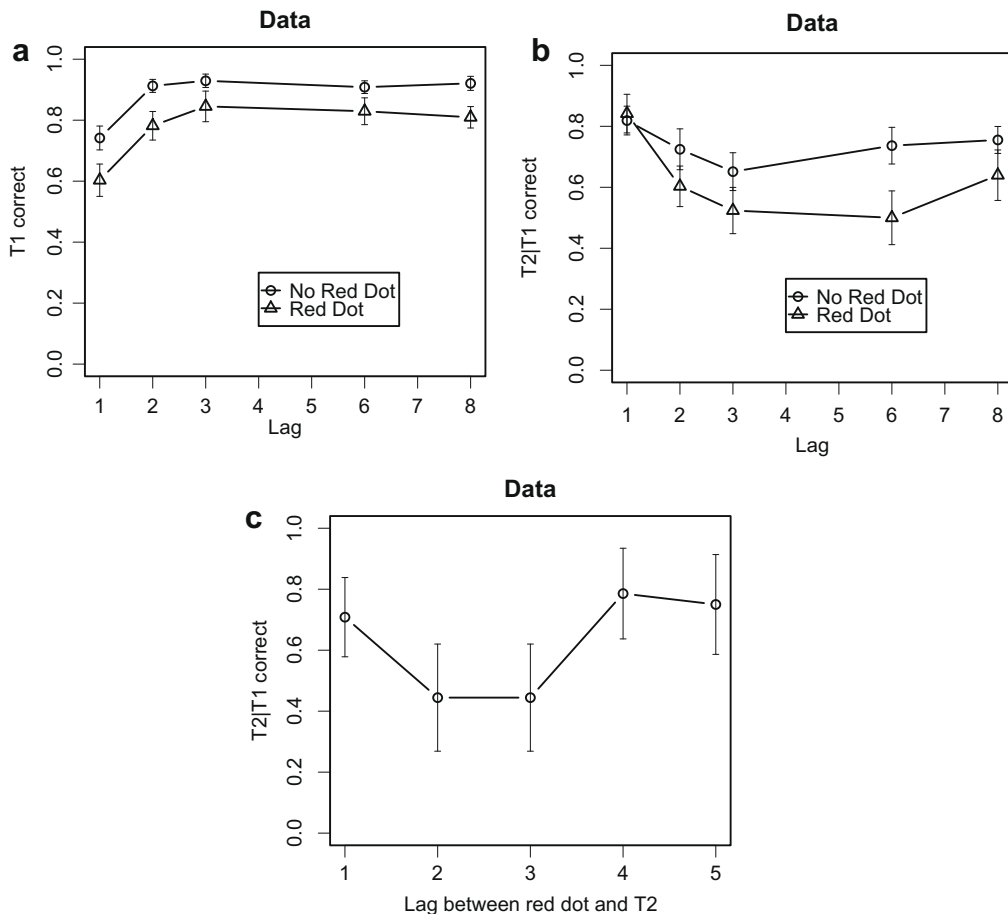


Fig. 14. The impact of a red dot on (a) T1 accuracy (b) T2|T1 accuracy, and (c) the impact of a red dot on T2|T1 accuracy for Lag-8, plotted with respect to the distance between the red dot and T2.

blink, because this only happens when there is enough time between T1 and T2. Fig. 14c shows T2|T1 accuracy on the Lag-8 trials plotted against the distance between the red dot and T2. Although this suggests that the red dot indeed produces a blink on T2, there is insufficient data to do a meaningful significance test.

5.3. Discussion and model

The results replicate the beneficial effect of the red dot manipulation, which according to our model is due to more relaxed control. Some of the experimental details were different between Experiment 1 and 2, but this only demonstrates the robustness of the effect. One difference between the two experiments is that the dot detection rate is much higher in Experiment 2. A possible explanation for this is that the presentation rate in Experiment 1 is 90 ms/item, while it is 100 ms/item in Experiment 2. Although that is only a small difference, a speculative explanation is that the model can activate exactly two productions in 100 ms, but not in 90 ms, and therefore misses many dot-production cycles. The 50 ms production activation time has never been investigated in great detail, but it has historically been a value that has worked for many different models (Anderson, 2007).

The results also confirm our hypothesis that the introduction of a task element that uses the resource that is also used for memory consolidation increases the AB. Finally, we see that there is indeed a learning effect on top of the effects of task manipulation. All of this is well in line with the model. In order to fit the data for the rotated condition, we increased the memory consolidation time by 50 ms in order to model that identification of the rotated letter takes longer. We decreased the time to recognize that the first target is indeed a target, because the rotation itself (instead of retrieval of the category) identifies an item as a target. Fig. 15 shows the model fits.

In both the model and the data, the dot task results in a strongly reduced AB, rather than a total lack of the AB (as in some non-blinkers). This is because in the model the control production that protects consolidation is still there: it just has more competition. Especially in Lag-3 it sometimes has the opportunity to stop target detection and produce a blink. We did not attempt to model the slight decrement in performance on T1 performance in the rotated task, which we assume is due to misidentification of the rotated letter.

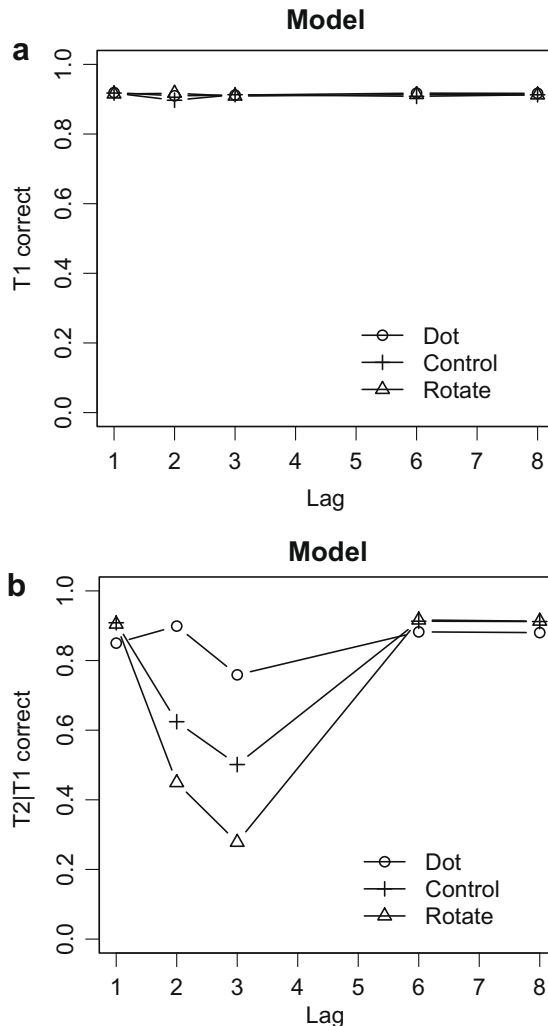


Fig. 15. Model fit of Experiment 2 results.

The model also reproduces the impact of the red dot on performance. The model's accuracy of detecting the dot is 81%, which is fairly consistent with the high, but not perfect performance in Experiment 2. Fig. 16a and b shows a comparison of the T1 and T2|T1 accuracies contrasting trials in which the dot turned red, and trials in which the dot stayed gray. Consistent with the data, the red dot produces an overall decrement in performance on T1 and a decrement in performance on T2|T1 for the longer lags. Finally, the model produces an AB on T2 caused by the red dot (Fig. 16c).

6. General discussion

The threaded cognition model we have presented in this article attributes the AB to cognitive control, consistent with recent models of the AB such as the temporary loss of control (Di Lollo et al., 2005) and the Boost and Bounce model (Olivers & Meeter, 2008). The main difference with these models is that it attributes the AB to an overexertion of control using a control production rule that is separate from both target detection and memory consolidation. This control production suppresses target detection for the duration of memory consolidation. The separation of control from other task knowl-

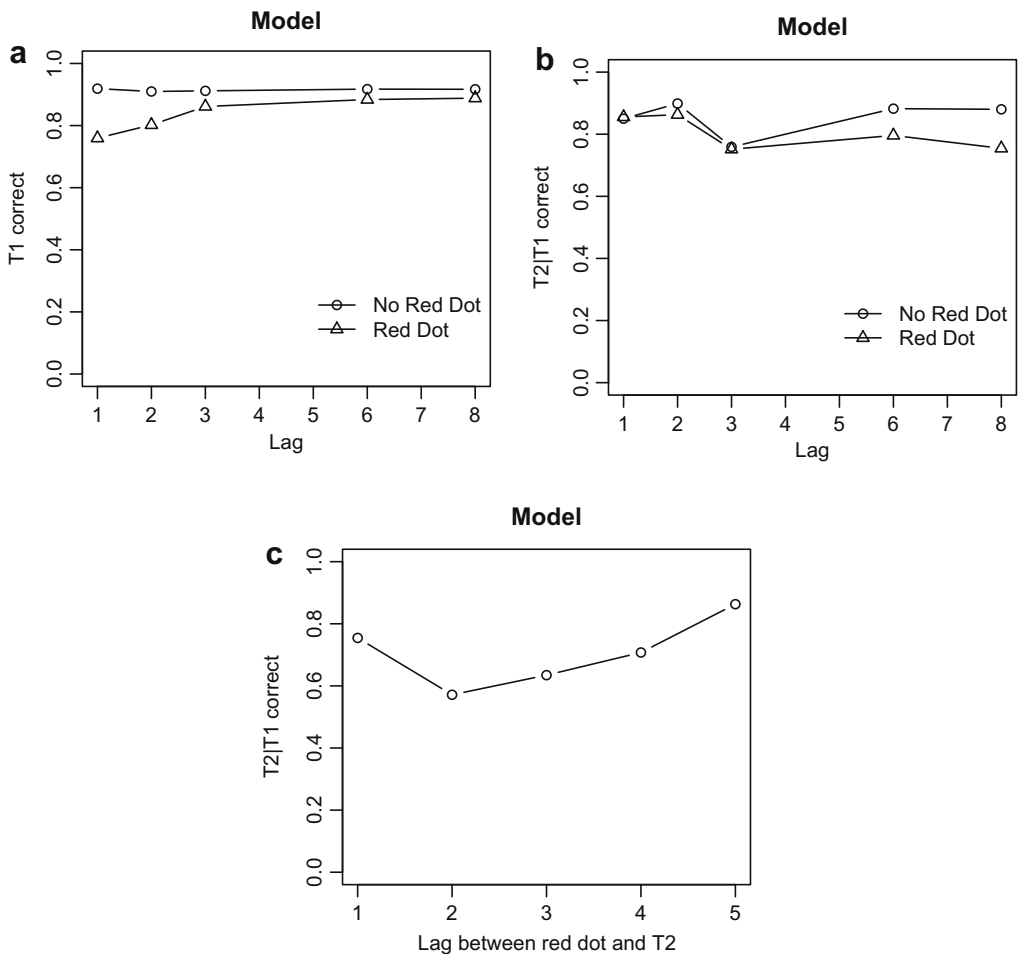


Fig. 16. Model fit of red dot impact: (a) T1 accuracy (b) T2|T1 accuracy, and (c) the impact of a red dot on T2|T1 accuracy for Lag-8, plotted with respect to the distance between the red dot and T2.

edge allows the model to capture situations in which a de-emphasis of the RSVP task, either by distraction or a secondary task, decreases the AB. As a critical test of this aspect of the model, we conducted two experiments that showed how the exertion of cognitive control is influenced by adding or removing emphasis on the RSVP task. A crucial assumption in the model, which is part of the threaded cognition theory, is that control is not exerted by a separate control system, but has to compete with task execution itself. This means that certain additional task demands, like adding the dot task, compete with control and therefore diminish the blink. The model also shares some characteristics with models that attribute the AB to the slow and serial nature of memory consolidation, in that the time needed for memory consolidation controls the period of time that target detection is suspended. In that sense to model is related to ST² (Bowman & Wyble, 2007), CODAM (Fragopanagos et al., 2005), and the original Chun and Potter (1995) model. As a consequence, the AB is tied to time, and not to the number of items.

The notion that the control aspect of the task should be considered separately from both target detection and memory consolidation is supported by a recent study by Colzato et al. (2008). This study contrasted bilinguals with monolinguals, showing that on a series of tasks bilinguals exert stronger inhibitory cognitive control. Inhibitory control is needed to suppress one language while speaking or processing the other, and the study shows this generalizes to non-linguistic tasks. One of the tasks was an RSVP task, and it was found that bilinguals had a stronger AB than monolinguals. This effect can be attributed to neither a target-detection nor a memory-consolidation process, because neither plays a role in speaking one as opposed to two languages.

A recent challenge to models that attribute the blink to control, or at least that the blink is triggered by a distractor, is an experiment by Nieuwenstein, Potter, and Theeuwes (in press). They expanded the Chun and Potter (1995) experiment where T1 was followed by a blank. The original study shows a decrease in blink, but the new study, in which T1 and T2 are separated by a variable number of blanks instead of distractors, shows an increase in blink when T2 is not presented for the standard 100 ms, but instead for only 58 ms followed by a 42 ms mask. The threaded cognition model does not cover this result straight away. However, because the model attributes the blink to control, a different task may lead to a different control strategy. In the Nieuwenstein et al. study, the goal in the trials with a variable number of blanks is to “catch” a T2 that is displayed very briefly after a sequence of blanks, instead of finding a target among distractors. This may well produce a different control strategy.

This also gives rise to a different issue: why do people insert control in the RSVP task? The dot manipulation shows that relaxing control is not counterproductive. Wyble, Bowman, and Nieuwenstein (in press) suggest that if memory consolidation is not protected, order information of targets is not properly maintained. As a consequence, targets are reported correctly, but in the wrong order. The insertion of a control production can also be prompted by instructions, in that the natural interpretation is to put target detection and memory consolidation in a sequence, instead of letting them operate in parallel. This can explain why different task instructions can lead to a decreased blink (Oliviers & Nieuwenhuis, 2006).

An important aspect of the model is that it is grounded in both the ACT-R (Anderson, 2007) and threaded cognition (Salvucci & Taatgen, 2008) theories, which are theories of cognition in general and of multi-tasking, respectively. As a consequence, the mechanisms that underlie the critical parts of the model are not designed to explain the AB. Indeed, the model itself is very simple (nine production rules, including the rules needed for the dot task), which is appropriate for a simple task. Additional factors that are known to influence the AB, such as bottom-up effects of highly salient items or manipulations of stimulus discriminability, remain to be accommodated into future versions of the model.

However, using only a limited number of parameters and assumptions, we have shown that the model accounts for a wide range of effects associated with the AB phenomenon. To the best of our knowledge, it is also the first computational model that provides an explicit account of target selection processes in non-blinkers. Rather than a lack of attentional resources or a temporary loss of control, we suggest that the AB reflects the exertion of too much cognitive control. As demonstrated in both our simulations and behavioral experiments, the amount of control can be manipulated such that individuals temporarily change into either non-blinkers or strong blinkers.

Acknowledgments

This research was supported by Office of Naval Research Grant N00014-08-10541 and Air Force Office of Scientific Research Grant FA95500710359 awarded to N. Taatgen, and a VENI grant from the Netherlands Organization for Scientific Research awarded to S. Martens. We thank Seth Herd and David Jilk for their contributions to the Leabra model, Nikola Valchev for doing a pilot experiment, and Stefani Nellen for her comments on the manuscript.

Appendix A

In this appendix we will discuss the details of the model and the parameters that were used to fit the data. We will first discuss the production rules in the model, and then the parameters of the model.

A.1. Production rules

We will present the rules in pseudo-English format. Each rule matches information available from one or more of the cognitive modules: Visual, Visual-location, Declarative, Imaginal and Goal. In this model, Visual-location (visual-where in the main text) represents peripheral vision and is used to first focus on the fixation cross, and in the dot task to track the color of the moving dot. Visual (visual-what) is used to perceive the items in the RSVP stream. Declarative is used to determine the category of the perceived item (letter or digit). Imaginal is used to consolidate targets, and Goal is used to identify the task, to track the control state of the model, and to temporarily hold an item before it is consolidated in the imaginal buffer.

The first rule, attend-item, is used to direct visual attention to the fixation cross. It is only used at the start of a trial, and is absent in the model traces because it activates before the first item appears on the screen.

Attend-item:

```
IF goal: the task is an RSVP task
  AND visual-location: there us an unattended stimulus
  AND visual: visual attention is not focused on anything
THEN visual: move attention to the unattended stimulus
  AND goal: set the control state to detect
```

The second rule, potential-target, initiates the retrieval of the category of a perceived visual item from declarative memory.

Potential-target:

```
IF goal: the task is an RSVP task in state detect
  AND visual: an item has been perceived
  AND declarative: is not in use or a distractor has been retrieved
THEN declarative: retrieve the category of the perceived item
```

This rule only initiates a retrieval: the result of the retrieval is handled by the third rule:

Item-isa-target:

```
IF goal: the task is an RSVP task in state detect
  AND declarative: the retrieved item is of the target category (letter)
THEN goal: add the item to the goal
```

This rule is then followed by the rule that initiates memory consolidation:

Store-target:

IF goal: the task is an RSVP task in state detect and an item has been added

AND imaginal: the imaginal module is not busy

THEN imaginal: initiate memory consolidation of the item

AND goal: remove item from the goal

The Item-isa-target and Store-target rules are combined into a single rule that immediately consolidates a target if the imaginal system is available. However, if the imaginal system is not available when a target has been retrieved, the separate rules are used. ACT-R's production compilation mechanism (Taatgen & Anderson, 2002) would normally learn this rule automatically, but we have added it by hand for this model.

Combined-Store-target:

IF goal: the task is an RSVP task in state detect

AND declarative: the retrieved item is of the target category (letter)

AND imaginal: the imaginal module is not busy

THEN imaginal: initiate memory consolidation of the item

The model further contains the two control productions, one to protect consolidation, and one to go back to target detection once consolidation is done:

Protect-Consolidation:

IF goal: the task is an RSVP task in state detect

AND imaginal: the imaginal module is busy

AND declarative: an item category has been retrieved

THEN goal: set the state to consolidate

Done-Consolidation:

IF goal: the task is an RSVP task in state consolidate

AND imaginal: the imaginal module is not busy

THEN goal: set the state to detect

A.2. Adding the dot task

The dot task (Experiments 1 and 2) is modeled by two additional productions. The first production dismisses the gray dot.

Dismiss-grey-dot:

IF goal: the task is an RSVP task

AND visual-location: there is a grey object in peripheral vision

THEN visual-location: dismiss the object and clear the visual-location

ACT-R requires an explicit clear of the visual-location, otherwise it will not detect the next move of the dot. If a red object is detected, the following rule is activated:

Dot-is-target:

IF goal: the task is an RSVP task

AND visual-location: there is a red object in peripheral vision

THEN goal: add "red dot" as item to the goal

The store-target rule will subsequently consolidate the fact that the red dot has been detected.

A.3. Conflict resolution

If more than one production rule is applicable at the same time, the rule with the highest *utility* value is chosen. In this model, each rule is assigned a fixed utility value that determines the order of priority, which produces the following order:

1. Potential-target and Combined-store-target.
2. Item-isa-target and Store-target.
3. Dismiss-grey-dot and Dot-is-target.
4. Protect-Consolidation, Attend-item, and Done-Consolidation.

If more than one production rule on this list matches, the rule higher on the list is chosen. Rules on the same line in the list are never applicable at the same time, so they never cause a conflict. If a rule loses the competition, it can in theory still activate in the next cycle, provided its conditions still apply, and no other rule with a higher utility is applicable at that moment.

A.4. Parameters

The parameters of the model control the duration of processing in the different modules. The duration of activating a production rule is fixed at 50 ms, but other durations can vary:

- Visual attention latency: By default, it takes 85 ms for the visual module to recognize a visual object. For the model, we drew a recognition time from a uniform distribution between 46 and 105 ms. If the recognition time exceeds the presentation rate (usually 100 ms/item), the model does not recognize that particular object.
- Imaginal delay: By default, it takes 200 ms to consolidate an item in the imaginal buffer. For the model, we drew a consolidation time from a uniform distribution between 150 and 450 ms.
- Declarative retrieval time: Retrieval time depends on the activation of the item to be retrieved from memory, using the following equation: $\text{retrieval time} = Fe^{-A}$. In this equation, F is a scaling parameter that we estimated at 0.2, and A is the activation of the item in memory, which in this model can be simplified to: $A = B + \text{noise}$. B is the base-level activation that we estimated at 2.0, and the noise is drawn from a logistic distribution with $M = 0$ and $\sigma^2 = 0.132$. These equations and parameter values result in an average retrieval time of 27 ms. The F and noise parameters have no default values in ACT-R.

To fit the data from the various experiments, we used the following changes to parameters:

- Basic Blink, Fig. 3: no changes.
- Di Lollo et al. (2005), Fig. 4: no changes.
- Bowman and Wyble (2007), Fig. 5: no changes in the 100 ms presentation rate condition. For the 50 ms presentation rate, the visual attention delay was drawn from 30 to 67 ms (resulting in anything with a delay >50 ms to be not recognized), and the Protect-Consolidation rule was given the same utility as the Potential-target and Combined-store-target rules. In case of a conflict between those rules, one was randomly chosen.
- Chun and Potter (1995), Fig. 6: The F parameter was set to 0.1.
- Experiment 2, Fig. 15: the Imaginal Delay was drawn from a uniform distribution between 150 and 350 ms. Rotated items received a higher activation in memory ($B = 3$), and the Imaginal Delay in the Rotated condition was drawn from a uniform distribution between 200 and 400 ms.

The full model is available through <http://act-r.psy.cmu.edu/models>.

References

- Alivisatos, B., & Petrides, M. (1997). Functional activation of the human brain during mental rotation. *Neuropsychologia*, *35*, 111–118.
- Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* New York: Oxford university press.
- Arend, I., Johnston, S., & Shapiro, I. (2006). Task-irrelevant visual motion and flicker attenuate the attentional blink. *Psychonomic Bulletin & Review*, *13*(4), 600–607.
- Bowman, H., & Wyble, B. (2007). The simultaneous type, serial token model of temporal attention and working memory. *Psychological Review*, *114*(1), 38–70.
- Chua, F. K., Goh, J., & Hon, N. (2001). Nature of codes extracted during the attentional blink. *Journal of Experimental Psychology-Human Perception and Performance*, *27*(5), 1229–1242.
- Chun, M. M., & Potter, M. C. (1995). A 2-Stage Model for Multiple-Target Detection in Rapid Serial Visual Presentation. *Journal of Experimental Psychology-Human Perception and Performance*, *21*(1), 109–127.
- Colzato, L. S., Bajo, M. T., van den Wildenberg, W., Paolieri, D., Nieuwenhuis, S., La Heij, W., et al (2008). How does bilingualism improve executive control? A comparison of active and reactive inhibition mechanisms. *Journal of Experimental Psychology-Learning Memory and Cognition*, *34*(2), 302–312.
- Colzato, L. S., Spapé, M. M. A., Pannebakker, M. M., & Hommel, B. (2007). Working memory and the attentional blink: Blink size is predicted by individual differences in operation span. *Psychonomic Bulletin & Review*, *14*(6), 1051–1057.
- Craston, P., Wyble, B., Chennu, S., & Bowman, H. (in press). The attentional blink reveals serial working memory encoding: Evidence from virtual & human event-related potentials. *Journal of Cognitive Neuroscience*.
- Dehaene, S., Sergent, C., & Changeux, J. P. (2003). A neuronal network model linking subjective reports and objective physiological data during conscious perception. *Proceedings of the National Academy of Sciences of the United States of America*, *100*(14), 8520–8525.
- Di Lollo, V., Kawahara, J., Ghorashi, S. M. S., & Enns, J. T. (2005). The attentional blink: Resource depletion or temporary loss of control? *Psychological Research*, *69*(3), 191–200.
- Dickerson, J., & Humphreys, G. W. (1999). On the identification of misoriented objects: Effects of task and level of stimulus description. *European Journal of Cognitive Psychology*, *11*(2), 145–166.
- Feinstein, J. S., Stein, M. B., Castillo, G. N., & Paulus, M. P. (2004). From sensory processes to conscious perception. *Consciousness and Cognition*, *13*(2), 323–335.
- Fragopanagos, N., Kockelkoren, S., & Taylor, J. G. (2005). A neurodynamic model of the attentional blink. *Cognitive Brain Research*, *24*(3), 568–586.
- 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.
- Heil, M. (2002). The functional significance of ERP effects during mental rotation. *Psychophysiology*, *39*, 535–545.
- 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.
- Leek, E. C., & Johnstons, S. J. (2006). A polarity effect in misoriented object recognition: The role of polar features in the computation of orientation-invariant shape representations. *Visual Cognition*, *13*(5), 573–600.
- Martens, S., Johnson, A., Bolle, M., & Borst, J. (2009). A quick visual mind can be a slow auditory mind: Individual differences in attentional selection across modalities. *Experimental Psychology*, *56*(1), 33–40.
- Martens, S., Munneke, J., Smid, H., & Johnson, A. (2006). Quick minds don't blink: Electrophysiological correlates of individual differences in attentional selection. *Journal of Cognitive Neuroscience*, *18*(9), 1423–1438.
- Martens, S., & Valchev, N. (2009). Individual differences in the attentional blink: The important role of irrelevant information. *Experimental Psychology*, *56*(1), 18–26.
- Nieuwenhuis, S., Gilzenrat, M. S., Holmes, B. D., & Cohen, J. D. (2005). The role of the locus coeruleus in mediating the attentional blink: A neurocomputational theory. *Journal of Experimental Psychology-General*, *134*(3), 291–307.
- Nieuwenstein, M. R., & Potter, M. C. (2006). Temporal limits of selection and memory encoding. *Psychological Science*, *17*(6), 471–475.
- Nieuwenstein, M. R., Potter, M. C., & Theeuwes, J. (in press). Unmasking the attentional blink. *Journal of Experimental Psychology-Human Perception and Performance*.
- Olivers, C. N. L., & Meeter, M. (2008). A boost and bounce theory of temporal attention. *Psychological Review*, *115*(4), 836–863.
- Olivers, C. N. L., & Nieuwenhuis, S. (2005). The beneficial effect of concurrent task-irrelevant mental activity on temporal attention. *Psychological Science*, *16*(4), 265–269.
- Olivers, C. N. L., & Nieuwenhuis, S. (2006). The beneficial effects of additional task load, positive affect, and instruction on the attentional blink. *Journal of Experimental Psychology-Human Perception and Performance*, *32*(2), 364–379.
- Olivers, C. N. L., van der Stigchel, S., & Hulleman, J. (2007). Spreading the sparing: against a limited-capacity account of the attentional blink. *Psychological Research*, *71*(2), 126–139.
- O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive Neuroscience*. Cambridge, MA: MIT Press.
- 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.
- Salvucci, D. D., & Taatgen, N. A. (2008). Threaded cognition: An integrated theory of concurrent multitasking. *Psychological Review*, *115*(1), 101–130.
- Schneider, W., Eschman, A., & Zuccolotto, A. (2002). *E-Prime user's guide*. Pittsburgh: Psychology Software Tools.
- 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., & 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.
- Shih, S. I. (2008). The attention cascade model and attentional blink. *Cognitive Psychology*, *56*(3), 210–236.

- Taatgen, N. A. (2005). Modeling parallelization and flexibility improvements in skill acquisition: From dual tasks to complex dynamic skills. *Cognitive Science*, 29(3), 421–455.
- Taatgen, N. A., & Anderson, J. R. (2002). Why do children learn to say “Broke”? A model of learning the past tense without feedback. *Cognition*, 86(2), 123–155.
- Taatgen, N. A., Juvina, I., Herd, S., Jilk, D., & Martens, S. (2007). Attentional Blink: An internal traffic jam? In: *Proceedings of the eighth international conference on cognitive modeling* (pp. 91–96). New York: Psychology Press.
- Wyble, B., Bowman, H., & Nieuwenstein, M. (in press). The attentional blink provides episodic distinctiveness: Sparing at a cost. *Journal of Experimental Psychology-Human Perception and Performance*.