



# Kent Academic Repository

**Bowman, Howard, Wyble, Brad, Chennu, Srivas and Craston, Patrick (2008)**  
***A Reciprocal Relationship Between Bottom-up Trace Strength and the Attentional Blink Bottleneck: Relating the LC-NE and ST2 Models.*** *Brain Research*,  
1202 . pp. 25-42. ISSN 0006-8993.

## Downloaded from

<https://kar.kent.ac.uk/23957/> The University of Kent's Academic Repository KAR

## The version of record is available from

<https://doi.org/10.1016/j.brainres.2007.06.035>

## This document version

UNSPECIFIED

## DOI for this version

## Licence for this version

UNSPECIFIED

## Additional information

## Versions of research works

### Versions of Record

If this version is the version of record, it is the same as the published version available on the publisher's web site. Cite as the published version.

### Author Accepted Manuscripts

If this document is identified as the Author Accepted Manuscript it is the version after peer review but before type setting, copy editing or publisher branding. Cite as Surname, Initial. (Year) 'Title of article'. To be published in *Title of Journal* , Volume and issue numbers [peer-reviewed accepted version]. Available at: DOI or URL (Accessed: date).

## Enquiries

If you have questions about this document contact [ResearchSupport@kent.ac.uk](mailto:ResearchSupport@kent.ac.uk). Please include the URL of the record in KAR. If you believe that your, or a third party's rights have been compromised through this document please see our [Take Down policy](https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies) (available from <https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies>).

available at [www.sciencedirect.com](http://www.sciencedirect.com)[www.elsevier.com/locate/brainres](http://www.elsevier.com/locate/brainres)


---



---

**BRAIN  
RESEARCH**


---



---



---

**Research Report**

# A reciprocal relationship between bottom-up trace strength and the attentional blink bottleneck: Relating the LC–NE and ST<sup>2</sup> models

H. Bowman, B. Wyble\*, S. Chennu, P. Craston

Centre for Cognitive Neuroscience and Cognitive Systems, University of Kent, Canterbury, Kent, CT2 7NF, UK

---

**ARTICLE INFO**
**Article history:**

Accepted 19 June 2007

Available online 30 June 2007

**Keywords:**

Temporal attention

Neural modeling

Attentional blink

Locus coeruleus

**ABSTRACT**

There is considerable current interest in neural modeling of the attentional blink phenomenon. Two prominent models of this task are the Simultaneous Type Serial Token (ST<sup>2</sup>) model and the Locus Coeruleus–Norepinephrine (LC–NE) model. The former of these generates a broad spectrum of behavioral data, while the latter provides a neurophysiologically detailed account. This paper explores the relationship between these two approaches. Specifically, we consider the spectrum of empirical phenomena that the two models generate, particularly emphasizing the need to generate a reciprocal relationship between bottom-up trace strength and the blink bottleneck. Then we discuss the implications of using ST<sup>2</sup> token mechanisms in the LC–NE setting.

© 2007 Elsevier B.V. All rights reserved.

---



---

**1. Introduction**

The attentional blink (AB) (Chun and Potter, 1995; Raymond et al., 1992) is one of the most studied temporal attention tasks. It is of particular interest since, firstly, it elucidates the length of an attentional episode (Bowman and Wyble, 2007) and, secondly, it has been shown to reflect a very late attentional bottleneck (Luck et al., 1996; Rolke et al., 2001; Shapiro et al., 1997a,b; Vogel et al., 1998). The latter of these makes it a particularly interesting task for studying the nature of conscious perception (Dehaene et al., 2006). Specifically, targets that are missed during the blink are nonetheless processed for sensory and semantic features (Luck et al., 1996; Rolke et al., 2001; Shapiro et al., 1997a,b; Vogel et al., 1998). Thus, it seems that blinked targets are extensively evaluated;

however, they do not quite receive sufficient processing to reach conscious awareness (Dehaene et al., 2006). Arguably then, the AB isolates the essential aspects of conscious perception; that is, the difference between missed and seen T2s exactly reflects the additional processes that are required to turn a rich preconscious representation into a conscious percept (Dehaene et al., 2006). In a similar vein, the AB also seems to be revealing with regard to the nature of working memory (WM), viz, the blink can be argued to reflect the essential processes that are required for an item to be encoded into working memory (WM).

The majority of attentional blink (AB) research has been behavioral in nature. However, there has been considerable recent interest in (1) identifying neural correlates (Marois et al., 2000; Martens et al., 2006; Nieuwenhuis et al., 2005b; Rolke

---

\* Corresponding author. Fax: +44 1227 762811.

E-mail addresses: [H.Bowman@kent.ac.uk](mailto:H.Bowman@kent.ac.uk) (H. Bowman), [bwbyble@gmail.com](mailto:bwbyble@gmail.com) (B. Wyble).

URL: <http://www.cs.kent.ac.uk/people/staff/hb5> (H. Bowman).

Abbreviations: ST<sup>2</sup>, Simultaneous Type Serial Token model; LC, locus coeruleus; NE, norepinephrine; AB, attentional blink; WM, working memory; RSVP, rapid serial visual presentation; ERP, event related potential; TAE, transient attentional enhancement; TFL, task filtered layer; SOA, stimulus onset asynchrony; TPJ, temporo-parietal junction

et al., 2001; Vogel and Luck, 2002; Vogel et al., 1998) and (2) developing neurally explicit computational explanations (Battye, 2003; Bowman and Wyble, 2007; Chartier et al., 2004; Dehaene et al., 2003; Fragopanagos et al., 2005; Nieuwenhuis et al., 2005b). Amongst these, the Simultaneous Type Serial Token (ST<sup>2</sup>) (Bowman and Wyble, 2007) and the Locus Coeruleus–Norepinephrine (LC–NE) (Nieuwenhuis et al., 2005b) models are notable approaches, since the former reproduces a broad spectrum of AB phenomena and the latter ties into a concrete neurophysiological substrate. In a sense, the ST<sup>2</sup> model works down from the behavioral data, while the LC–NE model works up from the neurophysiology. Thus, an important research question concerns how these two models relate to one another. In particular, to what extent can the neurophysiological mechanisms highlighted in the LC–NE model be reconciled with the cognitive-level mechanisms inherent in the ST<sup>2</sup> model? This is the issue that we explore in this article.<sup>1</sup>

Although we will consider a number of empirical phenomena, a particular focus will be on the robust finding that the AB is attenuated when a blank is placed after the first target (T1) (Chun and Potter, 1995). This is of particular significance since it reflects a more general finding that the easier the T1 task, the shallower the blink (Bowman and Wyble, 2007; Chun and Potter, 1995; Jolicoeur, 1998; Wyble and Bowman, 2005). In neural network terms, this phenomenon suggests that the greater the T1 bottom-up trace strength, the smaller the AB. That is, that there is a reciprocal relationship between bottom-up trace strength and the AB bottleneck.

A similarity between the LC–NE and ST<sup>2</sup> approaches is that they both assume attentional enhancements that are triggered by detection of a salient stimulus (e.g. a target stimulus) and are short-lived. In ST<sup>2</sup>, the enhancement is the transient attentional enhancement (TAE), which is realized by a mechanism called the *blaster*, while in the LC–NE model the enhancement arises from the locus coeruleus itself. That said, the two enhancements do have somewhat different purposes. ST<sup>2</sup>'s blaster aids detection and WM encoding of weak items, while the LC is assumed to have its affect on decision and response systems, where it optimizes processing in multilayer decision systems (Aston-Jones and Cohen, 2005). However, both approaches assume that it is unavailability of this attentional resource (TAE-blaster and LC respectively) that is the direct cause of the blink. The mechanism that causes this unavailability is, though, very different in the two cases. In ST<sup>2</sup>, the TAE-blaster is held offline by ongoing Working Memory (WM) encoding. This preserves the integrity of such encoding by preventing a second item from corrupting the episodic integrity of the encoding of a first item. In the LC–NE model, after firing, the LC enters an intrinsic refractory period in which it is difficult to re-fire. It is this difference between the approaches that is at the heart of the comparison we make in this article.

More specifically, we will argue that the previously mentioned reciprocal relationship between bottom-up trace strength and the AB bottleneck suggests that unavailability

of the attentional enhancement should, indeed, be tied to working memory encoding. This is because more strongly represented items would be encoded faster into WM, causing the attentional enhancement to be offline for a shorter period, thereby alleviating the blink. We will also highlight a neural mechanism, which we call a gate–trace pair that ensures that more strongly represented items are indeed encoded more rapidly into WM.

We will begin (in Section 2) by introducing the AB and highlighting what we consider to be the key behavioral phenomena that computational models should replicate. Then we review the ST<sup>2</sup> and LC–NE models in Section 3. We follow this with an assessment of the two models in Section 4. Then we explore the addition of an ST<sup>2</sup>-style gate–trace mechanism to the LC–NE model in Section 5. Finally, we discuss the implications of our work in Section 6.

## 2. The attentional blink

### 2.1. The task

The AB is commonly observed during Rapid Serial Visual Presentation (RSVP) (e.g. Chun and Potter, 1995; Raymond et al., 1992), in which a sequence of items are presented at the same spatial location with each rapidly replacing the previous item; the rate of presentation typically being around 10 items per second (unless otherwise stated, we assume such a presentation rate throughout the article). At such speeds, items only yield fleeting mental representations and targets within the sequence can be missed. The AB (Raymond et al., 1992) can be observed in RSVP tasks in which two targets are placed in the sequence. If the first target (T1) is correctly reported, report of the second target (T2) is impaired when it appears within 200 to 500 ms of the onset of the first target.

In this paper, we focus on Chun and Potter's (1995) blink task in which participants must report the identity of two letter targets (T1 and T2) presented in a stream of digit distractors (we call this the *letters-in-digits* task). This task can be argued to yield a *pure* test of the blink, in the sense that there is no task switch between the T1 and the T2 tasks, which has been argued to be a potential confound (Chun and Potter, 2000).

### 2.2. The AB phenomena

We claim that the following are key AB phenomena that computational models should reproduce.

1. A basic blink curve with lag-1 sparing.
2. Increased processing of T1+1 slot.
3. Blink attenuation with T1+1 blank.
4. Blink attenuation with T2+1 blank.
5. Delayed consolidation for T2s seen during the blink.
6. Spatial specificity of the lag-1 enhancement.
7. T1–T2 costs at lag-1.
8. Late stage bottleneck.
9. Spreading the sparing.

We discuss these phenomena in turn.

<sup>1</sup> We concentrate on just two models here. However, Bowman and Wyble (2007) present a detailed review of the full spectrum of models.

### 2.2.1. The basic blink

A typical AB serial-position curve, arising from the letters-in-digits task (Chun and Potter, 1995) is shown as the Basic Blink condition in Fig. 3(b). Points to note are,

1. the blink is a 200 ms–500 ms (approx.) interval post T1 onset in which performance on T2, conditional on correct report of T1 (i.e. T2/T1) is significantly reduced;
2. generally the blink has a sharper onset than offset; and
3. if T2 immediately follows T1 it is reported at baseline levels, which is described as *lag-1 sparing*.

### 2.2.2. Increased processing of T1+1 slot

There is a good deal of evidence that the item (whether it be a distractor or a target) immediately after the first target in a dual target RSVP stream is particularly deeply processed. For example, in a letter detection AB paradigm, Chua et al. (2001) found that a distractor immediately following a T1 primes a later T2 more than it would at other positions relative to T1. This finding suggests that the T1 opens up a short window of enhancement, which includes the following distractor.

Furthermore, lag-1 sparing suggests increased processing when the T1+1 item is a target. In fact, in some blink tasks (in particular, those in which there is no task switch between T1 and T2) T2 performance at lag-1 can be above baseline; that is, above the post recovery (e.g. lags 7 and 8) level of T2 performance; see section 5.2 in Bowman and Wyble (2007).

Thus, it seems that the occurrence of the T1 initiates a brief window of generalized attentional enhancement. Furthermore, there is evidence that this window has a fixed minimal extent; that is, it lasts at least 120 ms (Potter et al., 2002) also, see section 5.1 of Bowman and Wyble (2007) and experiment 1 of Nieuwenhuis et al. (2005b). The emphasis here is on “minimal extent”, since there is evidence that the window can be extended when a sequence of target items is presented; see the Spreading in the Sparring section to follow.

### 2.2.3. Blink attenuation with T1+1 blank

The blink is attenuated if a blank is placed in the T1+1 position, but not if the blank is placed at T1+2 (Chun and Potter, 1995; Raymond et al., 1992), see Fig. 3(b). This suggests that when T1 is easier to perceive, T2 is also more easily perceived.

It is generally accepted (e.g. see Seiffert and Di Lollo, 1997) that, although weak in the sense of classic masking studies (e.g. Enns and Di Lollo, 2000), each RSVP item masks the item that precedes it. Thus, the activation of a node representing an item is curtailed by the arrival of the mask. Therefore, an item that is masked has a shorter activation trace than an unmasked item. Attenuation of the blink with T1+1 blank then suggests that unmasked targets are consolidated into WM more easily than masked targets. Although, see Olivers (2007) and Raymond et al. (1992) for an alternative perspective, which suggests that the blink is the result of a period of reactive inhibition, which is initiated to prevent the T1+1 distractor from entering WM. However, a comparison between the bottleneck models presented here and this theory awaits a concrete computational instantiation of the reactive inhibition account.

We regard blink attenuation with T1+1 blank to be a particularly significant phenomenon, since it reflects a broad finding that making T1 processing easier attenuates the blink (Bowman and Wyble, 2007; Chun and Potter, 1995; Jolicoeur, 1998; Wyble and Bowman, 2005).<sup>2</sup> In neural networks terms, this suggests that increased T1 activation attenuates the blink or, in other words, there is a reciprocal relationship between T1 bottom-up trace strength and the AB bottleneck.

### 2.2.4. Blink attenuation with T2+1 blank

In the same spirit, the strength of the T2 trace (as also regulated by backward masking) affects blink depth. Although empirical studies have not directly assessed this fact, it has been shown that the blink is absent if T2 is the last item in the stream (Giesbrecht and Di Lollo, 1998), where it will automatically be unmasked, see Fig. 3(b). This was confirmed by Sessa et al. (2006) and Vogel and Luck (2002). Thus, we will in fact model complete abolition of the blink with T2 as the last item in the stream.

In summary then, ease of target processing modulates blink depth. In fact, for both targets, there is a reciprocal relationship between bottom-up trace strength and the blink bottleneck, with stronger targets generating weaker blinks.

### 2.2.5. Delayed T2 consolidation

In typical AB studies, the blink is not total; that is, T2 performance is never zero at any lag. This raises the question of the fate of T2s seen during the blink. There are two extreme positions, viz that seen T2s *breakthrough* or *outlive* the blink. In particular, T2 manipulations that attenuate the blink (e.g. increasing the personal or emotional salience of the T2 (Anderson, 2005; Shapiro et al., 1997a)) are sometimes described as T2 breakthrough effects. However, it may be that this term is not completely appropriate. For example, improved T2 performance during the blink does not seem to reduce T1 identification, which one might expect if a T2 broke through, thereby causing it to compete with T1.

Furthermore, ERP studies suggest that T2 consolidation is delayed during the blink (Martens et al., 2006; Sessa et al., 2006; Vogel and Luck, 2002). The key measure considered being T2–P3 latency, i.e. the time from T2 stimulus onset to corresponding P3 peak. This raises the possibility that particularly salient or perceptually strong T2s survive the blink because they generate stronger representational traces that can outlive the blink. Conclusively distinguishing between breaking through and outliving is difficult; however, the finding of delayed T2 consolidation during the blink is revealing and serves as a clear cut finding that models should replicate.

### 2.2.6. Spatial specificity of lag-1 enhancement

As previously discussed, the lag-1 attentional enhancement is generalized, in the sense that an enhancement is observed whatever the lag-1 item (e.g. target or distractor). However, there is evidence that the enhancement is not spatially

<sup>2</sup> Although see McLaughlin et al. (2001), who do not find blink attenuation with T1 difficulty. However, we suspect that this finding is a result of the non-standard (skeletal) presentation employed and the brief targets and target-mask onset asynchronies employed.

generalized. In particular, Visser et al. (1999) have shown that there is no sparing if a lag-1 T2 appears in a different spatial location to T1, suggesting that the enhancement is restricted to the location of the initiating stimulus. This finding has been generalized to a spatial cueing setting (Wyble et al., submitted for publication-b).

### 2.2.7. T1–T2 costs at lag-1

Lag-1 sparing does not come free of cost. Initial evidence for this perspective is the reduced T1 performance at lag-1, see Fig. 3(d), suggesting competition between T1 and T2 at this lag (see also Potter et al., 2002). Further evidence of lag-1 costs arises from data on temporal order confusion; that is, situations in which T1 and T2 are both identified, but are “perceived” in the wrong order. As shown in Fig. 3(d), at lag-1, participants are only about 70% accurate at reporting the temporal order of targets (Chun and Potter, 1995). This deficit in order report disappears rapidly as the two targets are moved apart, reaching 95% accuracy by lag-3.

In fact, Bowman and Wyble (2007) go even further, claiming that lag-1 sparing is only observable for simple targets. They suggested that with more complex target types, for example, those with migratable component parts, lag-1 is not spared (see Bowman and Wyble, 2007, section 5.3). This was argued to reflect a loss of episodic distinctiveness, which, depending upon the form of the task employed, could arise as feature or component conjunction errors (Chun, 1997a), repetition blindness (Chun, 1997b) or order inversions (Chun and Potter, 1995).

As a reflection of this evidence, in the context of this paper and letters-in-digits tasks, we suggest that computational models should generate a decline in T1 performance and an increase in temporal order confusion at lag-1.

### 2.2.8. Late stage bottleneck

Although not per se a behavioral phenomenon, an important constraint on models is the late stage nature of the AB bottleneck. It is beyond the scope of this article to fully elaborate all the supporting data (see Bowman and Wyble, 2007, section 3.2 for a more detailed account), however, the data is extensive, including priming studies (Shapiro et al., 1997b), “breakthrough” effects (Anderson and Phelps, 2001; Shapiro et al., 1997a) and ERP data (Luck et al., 1996; Rolke et al., 2001; Vogel et al., 1998). All these studies suggest that the AB arises after sensory and also some degree of semantic processing. That is, even though a T2 may be missed during the blink, it is nonetheless extensively processed, in respect of the extraction of both visual and semantic features. ERP findings are particularly compelling in this respect since they have shown that early visual components (the N1 and P1) and semantic effects (the N400) are present irrespective of whether T2 is missed or seen. In contrast, a component typically associated with WM update (the P3) is reduced when a T2 is presented during the blink compared with when it is presented outside that period (Luck et al., 1996; Vogel et al., 1998).

### 2.2.9. Spreading the sparing

There is recent evidence that the blink is not absolute, in the sense that sparing can be extended beyond lag-1 if a continuous stream of targets is presented (Di Lollo et al.,

2005; Olivers et al., 2005). Although spreading the sparing is a newly discovered phenomenon, there are earlier findings in the literature that might be viewed as related to it. For example, it has often been observed that performance is spared at lag-2 in the T1+1 blank condition, see Fig. 3(b). However, (as justified in Section 4.2.3) the sequence T1 T2 (i.e. T1 blank T2) could be argued to be almost identical to T1 T1 T2 (since absence of a backward mask enables a target’s iconic representation to continue unabated). Thus, the sparing at lag-2 in the T1+1 blank condition could also be argued to arise because of the presence of a continuous stream of targets.

## 3. The models

We introduce the ST<sup>2</sup> and LC–NE models in turn, starting with ST<sup>2</sup>.

### 3.1. The ST<sup>2</sup> model

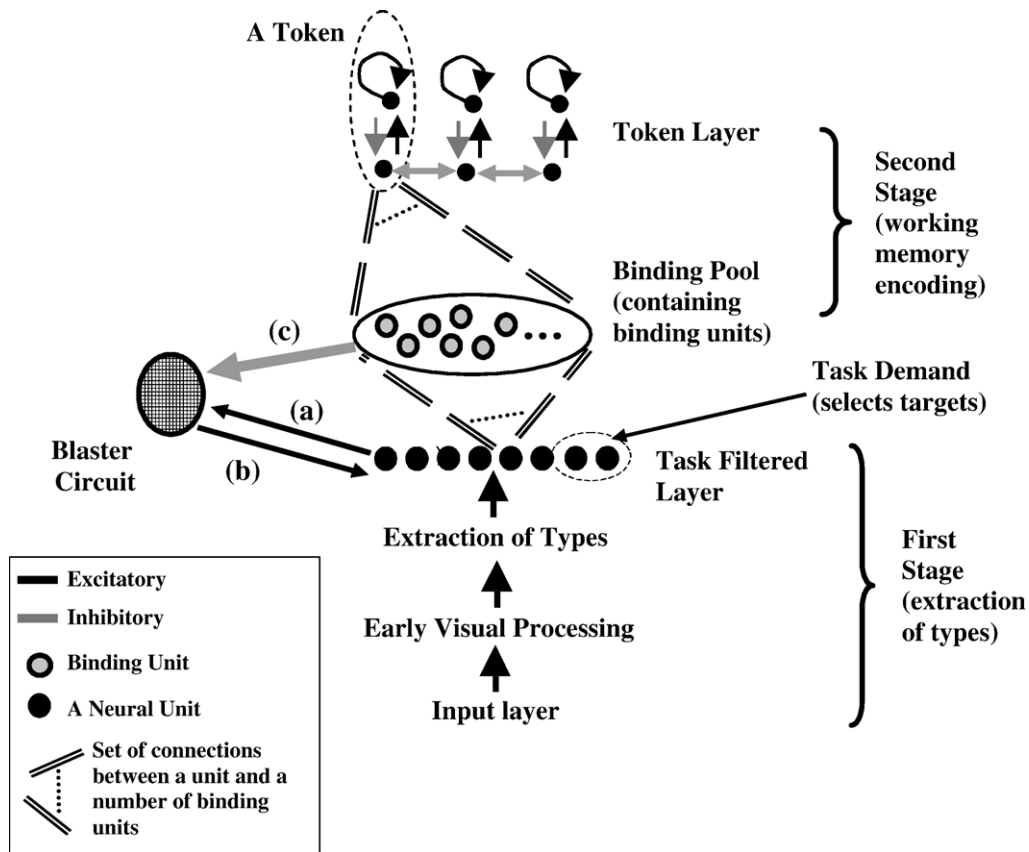
The Simultaneous Type, Serial Token model of temporal attention provides an abstract neural model of episodic attentional processing. The model incorporates constituent processes of such episodes: visual processing, item identification, attentional selection and encoding into working memory. ST<sup>2</sup> is targeted at modeling the letters-in-digits task (Chun and Potter, 1995), although ongoing work is seeking to expand the applicability of the model by considering tasks that contain a task switch and identify targets with colour marking. The model particularly encapsulates the *episodic distinctiveness hypothesis* of the AB (Bowman and Wyble, 2007); that is, that the AB reflects a system attempting to allocate unique episodic contexts to targets.

The overall architecture of the ST<sup>2</sup> model (strictly Neural-ST<sup>2</sup>) is shown in Fig. 1. Before discussing the components of the architecture, we highlight a central distinction used in the model.

#### 3.1.1. Types–tokens

The types–tokens distinction has been considered in the context of a number of temporal attention tasks (Chun, 1997b; Kanwisher, 1987, 1991). We use the term type to describe all featural properties associated with an item. This includes both sensory and semantic features. For example, the type representation of the letter K would contain, 1) all semantic features, e.g. that the item is in the category of letters and it follows L in the alphabet and 2) all the constituent features that comprise the visual representation of the letter, e.g. its shape, the angled line segments that comprise it and the colour in which it appears.

In contrast, a token represents instance specific (or episodic) information about the occurrence of an item. Thus, a token indicates that a particular type has occurred and also, when, relative in time to other items, it occurred. In the ST<sup>2</sup> model, WM encoding is the process of associating (or binding) a token to a currently active type; we also use the term *tokenization* to describe this process. In this sense, once bound, tokens act as “pointers”, from which the corresponding type can be regenerated when required, e.g. at retrieval.



**Fig. 1 – Structure of ST<sup>2</sup>.** Connectivity into and out of the binding pool is only shown schematically and for one type–token association, but applies for all type–token associations. This connectivity ensures that each token activates a portion of the binding pool, as does each type. The overlap between these portions defines a particular type–token association, for full details see [Bowman and Wyble \(2007\)](#). Only pairwise lateral inhibition is shown between token gates, but is present between all token gates. Conceptually, the binding pool and the blaster sit between stages 1 and 2. In particular, projections to and from the blaster arise from both stage 1 and the binding pool (which is post stage 1).

### 3.1.2. First stage

Input activation propagates through layers that abstractly represent steps of visual processing. These layers contain inhibitory interactions that generate masking effects. The fleetingness of item representations arises from these interactions, which determine bottom-up trace strength. Thus, activation traces are greatly weaker for a stimulus followed by another stimulus, as opposed to a stimulus followed by a blank; see ([Bowman and Wyble, 2007](#)). We differentiate between such traces, by describing them as either *strong* (followed by a blank) or *weak* (followed by a stimulus).

The early visual processing layers feed into layers that abstractly represent type oriented processing, including extraction of semantic and conceptual features ([Potter, 1993](#)). The “output” of stage 1 is the Task Filtered Layer (TFL), at which task demand foregrounds nodes corresponding to targets. The TFL is the gateway through which items (corresponding to fully processed types) can be encoded into tokens.

### 3.1.3. Second stage

The first stage yields a decaying trace of the visual and semantic features of target items within the TFL. The first

stage is parallel in the sense that different items can be simultaneously active at the TFL. In contrast, stage two implements sequential encoding of these items into WM, with this sequentiality emerging from mechanisms that attempt to ensure that items are discretely bound into WM. Before discussing the details of ST<sup>2</sup>'s token system, we highlight a general neural mechanism (called a gate–trace pair), of which ST<sup>2</sup> makes liberal use.

**3.1.3.1. Gate–trace active maintenance.** We introduce this mechanism in general terms (i.e. beyond the specifics of ST<sup>2</sup>), since we believe it tackles a problem in active memory in general. A standard way to set-up an activation-based WM ([Cowan, 2001](#)) is to include a layer of units with strong self-loops, with one unit per item that can be encoded, e.g. [Davelaar and Usher \(2004\)](#). Although its simplicity gives it value, there are difficulties with such an approach, see sections 2.5 and 6.1 of [Bowman and Wyble \(2007\)](#). Here we focus on one particular difficulty.

Specifically, such an active maintenance layer would not be well-behaved in a heavily competitive setting, i.e. in the presence of strong lateral inhibition. The idea that representations compete for access to attentional and perceptual

processes and indeed WM is well accepted (Desimone and Duncan, 1995). The difficulty is that, once encoded, i.e. in a self-sustaining attractor state, an item would impose strong competitive pressure on the remainder of the layer, potentially preventing any other item from being represented. As a result, with strong (selective) competition, such a WM system would have a capacity of one.<sup>3</sup> Although certainly not the only solution to this problem, gate–trace pairs are, we would argue, an elegant approach.

A gate–trace pair is an inhibitory interneuron circuit, involving an excitatory gate node and a self-sustaining inhibitory trace node. Each token in Fig. 1 has this format, with the gate node shown at the bottom of the circuit and the trace node at the top. The gate node is so named because it acts as a gate to the trace node. That is, activation of the trace node (which is in a closed circuit with the gate) is completely caused by the gate.<sup>4</sup> The trace node is so named because, once activated above threshold, it will self-sustain; thus it maintains the trace of an item in WM. However, once over threshold, the trace node will also suppress the gate node. Importantly, when gate nodes are placed in an inhibitory competition (see the tokens layer of Fig. 1), a receptive competitive active memory is obtained. Because trace nodes maintain representations and gates are removed from the competition once their traces have been allocated, the layer can engage in a series of (competitive) encoding episodes.

Gate–trace circuits have the further benefit that encoding time is inversely related to activation strength. The time to encode a stimulus using a gate–trace pair is the interval from when the gate starts to be active to when its corresponding trace node crosses threshold. Importantly, this time will be shorter the stronger the extrinsic (bottom-up) activation of the gate.<sup>5</sup> Because of this property of gate–trace pairs, they naturally yield a reciprocal relationship between bottom-up trace strength and encoding time, which we will shortly argue underlies the depth of the AB in ST<sup>2</sup>.

**3.1.3.2. Token binding.** During presentation of an RSVP stream, above threshold target items at the TFL can activate a token, thereby initiating a tokenization process. The token system is a competitive gate–trace layer with additional mechanisms that ensure tokens become available in a predefined order. While a token is being bound, binding pool units situated between the token's gate node and the active TFL node are incrementally allocated. Importantly, these units also have a gate–trace format; thus, they also have the

characteristic that the rate of encoding is proportional to the strength of the TFL activation.

As required by the ST<sup>2</sup> theory, tokens are made available sequentially during encoding. The gate nodes of the tokens compete to become available at the beginning of a trial and again whenever a token has completed binding, which inactivates its gate node. Thus, ST<sup>2</sup> ascribes unique episodic contexts and thereby preserves order of encoding when items are presented slowly enough.

It is possible that binding pool nodes allocated from different tokens may project to the same TFL node. This corresponds to the situation in which multiple instances of the same type have been bound, i.e. a repetition. It is also possible for the system to reach states in which there are binding pool units allocated from one token to many TFL nodes. At rapid presentation rates, such as those found in AB experiments, a target may become strongly active at the TFL while a previous target is being tokenized. In this situation, both items can be bound into the same token. That is, both items are encoded, but order information is lost.

At retrieval, the model first locates bound tokens (as indicated by active trace nodes) and then identifies binding units within the binding pool allocated from each token. If binding units are allocated, the type (or types) “pointed to” by those units is retrieved. The recall phase also determines perceived order of encoding. If a token has been bound, which points to multiple types then the model is prone to confusing the order of presentation; see Bowman and Wyble (2007).

#### 3.1.4. Binding pool

Types and tokens are bound via a pattern of sustained activation in a pool of binding units; see Fig. 1. In particular, there is no synaptic change involved in this binding mechanism (Bowman and Wyble, 2007). Thus, the binding pool is an activation-based memory for type–token associations. The projections into and out of the pool are arranged such that each type–token association is represented by a portion of the pool. We have explored both localist and distributed binding pools. The former would contain a unique binding unit for each type–token association, while the latter overcomes the lack of scalability of the former through overlapping type–token associations (Bowman and Wyble, 2007).

#### 3.1.5. Transient attentional enhancement: the blaster

ST<sup>2</sup> suggests that when the visual system detects an item that may be task relevant, a transient attentional enhancement (TAE) occurs, which is directed at the location at which that item appears. For a weak (masked) item, the contribution of this enhancement is critical in enabling it to activate a token and thereby be encoded.

The TAE is implemented by a mechanism called the blaster; see Fig. 1 (full implementation details can be found in Bowman and Wyble, 2007). Above threshold activity in any node of the TFL (which task demand ensures will only happen for targets) excites the blaster through the projection marked (a) in Fig. 1. The blaster sends a powerful excitatory projection to type nodes in stage-1 (through the projections marked (b) in Fig. 1). This causes a general, but short lived, excitation of TFL nodes.

<sup>3</sup> Note that the emphasis on strong (selective) competition is important here, since weaker forms of competition can be reconciled with higher capacity WMs (e.g. Davelaar and Usher, 2004).

<sup>4</sup> In this sense, the approach has similarities to a set of other models in which WM update is gated (Braver and Cohen, 2000; Frank et al., 2001).

<sup>5</sup> Modulo the nonlinearity in activation dynamics, the time to threshold for the trace node is directly proportional to the area under the gate node's activation curve; and the higher the gate activation, the quicker the area builds up.

3.1.5.1. *Suppression of the TAE by ongoing tokenization.* The cause of the blink is inhibition through the projection marked (c) in Fig. 1. This link ensures that, while binding units are being allocated, the blaster is held offline. Importantly, this inhibition is not active while an existing encoding is being maintained. Since binding units are in fact gate–trace circuits, the gate is only active during an allocation period; it is the gate of each binding unit that has an inhibitory projection to the blaster; see Bowman and Wyble (2007).

Therefore, the general behavior of the blaster is a brief spike of excitation followed by a period of inactivity until the completion of the current token binding. This is the mechanism by which the model exhibits an AB: unavailability of the blaster protects the integrity of the T1 binding by limiting attentional resources that could cause binding intrusions. This is consistent with our central theoretical position that the blink is the marker of a system that is attempting to allocate distinct episodic contexts. Thus, it “deliberately” sacrifices T2s in order to ensure the episodic integrity of T1, i.e. that it is unambiguously tokenized.

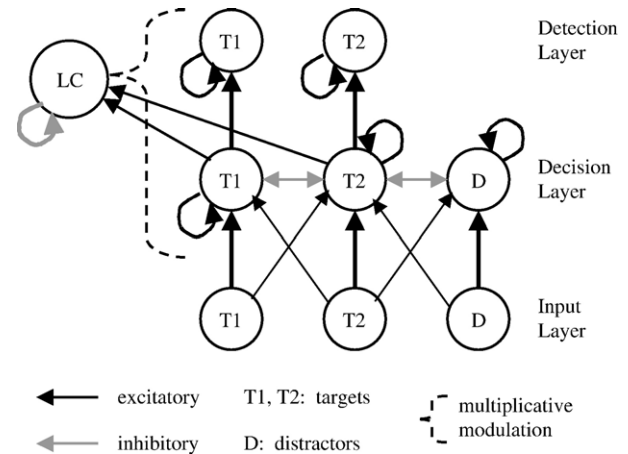
### 3.2. The LC–NE model

Nieuwenhuis et al.’s LC–NE model (Nieuwenhuis et al., 2005b) makes an important contribution to understanding the AB; a strength being that the model is framed within the context of a broad neurophysiological theory of attentional function (Aston-Jones and Cohen, 2005; Aston-Jones et al., 2000; Nieuwenhuis et al., 2005a; Usher et al., 1999). Perhaps of most importance, the LC–NE theory proposes a specific neurophysiological underpinning to the P3 (Nieuwenhuis et al., 2005a), which is being extensively used as an ERP correlate of attentional processing in the AB context (Kranzloch et al., 2003; Luck et al., 1996; Martens et al., 2006; Vogel and Luck, 2002; Vogel et al., 1998). Thus, the LC–NE theory proposes a unified explanation of temporal attention, which reconciles the AB phenomenon, neurophysiology, electrophysiology (in respect of the P3) and neural modeling.

The model explains the AB in terms of the functioning of the locus coeruleus (LC) (Nieuwenhuis et al., 2005b), a minute brain-stem structure (German et al., 1988) that projects widely to the cortex, with a special emphasis on areas involved in attentional processing (Aston-Jones et al., 2000). The model particularly focuses on the behavior of the LC while it is in phasic mode. In this mode of behavior, only salient items activate the LC, generating a phasic LC response. During such a response, LC innervation of the cortex ensures that the excitatory response of cortical neurons is amplified. In a visual discrimination task, monkey LC neurons activated with a temporal profile that seems to match the AB (Aston-Jones et al., 2000). This observation prompted the (Nieuwenhuis et al., 2005b) model; see Fig. 2. The model can be divided into two components: the behavioral network and the LC. We discuss these in turn, before considering how the model generates an AB.

#### 3.2.1. Behavioral network

The behavioral network is a simple feed-forward system, with major inter-layer connections being one-to-one. The network comprises three layers: *Input*, *Decision* and *Detection*; see Fig. 2. In the AB context, a sequence of stimuli is presented at the input layer to simulate the RSVP stream.



**Fig. 2** – This depiction of the LC–NE model is a redrawing of Fig. 2 of Nieuwenhuis et al. (2005b). Note, inhibitory and crosstalk connections between T1 and D are not shown for simplicity of presentation and point size of arrows indicates weight strength. This diagram is reproduced from Bowman and Wyble (2007).

The decision layer implements an ongoing competition between three alternatives: the two targets and a single unit abstractly modeling all distractors. Nodes in the decision layer compete through lateral inhibition. Crosstalk connections are also included between input and decision nodes, reflecting feature similarity between stimuli. The decision layer projects in one-to-one fashion to the detection layer, which serves as the output from the model. On the assumption that only targets are reported, the detection layer does not represent distractors. Finally, excitatory self-loops are included to sustain activation at decision and detection nodes. However, these loops are not strong enough to yield an active memory.

#### 3.2.2. LC firing

The LC circuit modulates activity in the behavioral network. Specifically, the LC is excited by detection of a salient stimulus. In the context of Nieuwenhuis et al.’s AB model, this means that targets fire the LC, as reflected by target nodes in the decision layer having excitatory projections to the LC. LC activity has a modulatory effect on the behavioral network, by simulating the release of Norepinephrine (NE). This release multiplicatively scales the afferent signals to network units, transiently adjusting their gain. Amplification of the multiplicative gain “sharpens” the sigmoidal transfer functions (Aston-Jones and Cohen, 2005) making decision and detection layer units temporarily more responsive.

#### 3.2.3. Refractory period—how the model blinks

However, after firing, the LC enters a refractory period. This arises since, while NE enhances processing in cortical areas, local NE release within the LC is believed to be autoinhibitory. Thus, following a phasic response, this autoinhibition generates an LC refractory period, during which further LC phasic response is rare (Nieuwenhuis et al., 2005b).

In the LC–NE model, it is unavailability of the LC during its refractory period that causes the blink. That is, the model



blinks as a result of the following sequence of events. Firstly, the T1 fires the LC, which provides a window of enhancement lasting around 150 ms. Secondly, following this T1-induced firing, the LC enters its refractory period. Thirdly, T2s arriving during this period are unable to immediately re-fire the LC and, consequently, do not benefit from a timely (NE-induced) increase in gain. This leaves T2s arriving during the refractory period at a disadvantage, ensuring that only particularly strong T2s get reported.

Blink recovery arises since, if the T2 follows the T1 after a sufficient time interval, the LC will have recovered from its refractory period and the T2 will be able to fire it. Indeed, the LC-NE model suggests that the AB profile will exactly follow the profile of the LC refractory period.

## 4. Assessment of models

We assess the two models against the empirical phenomena previously highlighted. We start with  $ST^2$ .

### 4.1. The $ST^2$ model

#### 4.1.1. The basic blink

The  $ST^2$  model generates a blink because the blaster is suppressed by ongoing T1 tokenization. T2s at lags 2 and 3

fall at the point of maximum impairment (see Fig. 3(a)), due to the length of time they have to wait for the blaster to come back on line. The impairment decreases monotonically through lags 4, 5 and 6, as it becomes more likely that tokenization of the T1 finishes before T2 has decayed.

However, at lag-1, T2 is close enough in time to T1 to take advantage of the (T1-initiated) blaster firing. Thus, the model demonstrates lag-1 sparing (see Fig. 3(a)). However, although T2 is typically encoded at lag-1, invariably, this only occurs into the first token, alongside T1.

#### 4.1.2. Increased processing of T1+1 slot

$ST^2$ 's TAE provides an enhancement that begins shortly after an initiating item (the T1 in an AB setting) and is very brief (lasting around 50 ms). Thus, in an AB setting, the blaster enhances the T1 and T1+1 slots. There are a number of phenomena that this mechanism enables the model to exhibit. Firstly, as just discussed, the model generates lag-1 sparing. Secondly, in fact the model's lag-1 performance has a tendency to be elevated above baseline (i.e. post recovery and single target performance); see Fig. 3(a) and the 50 ms SOA data in Fig. 19 of Bowman and Wyble (2007). (As discussed in Section 2.2.2, this is also found in humans.) Thirdly, in Bowman and Wyble (2007) we reproduced Chua et al.'s (2001) finding that a distractor is a more effective prime of a T2 if it is preceded by a T1. Finally, although we

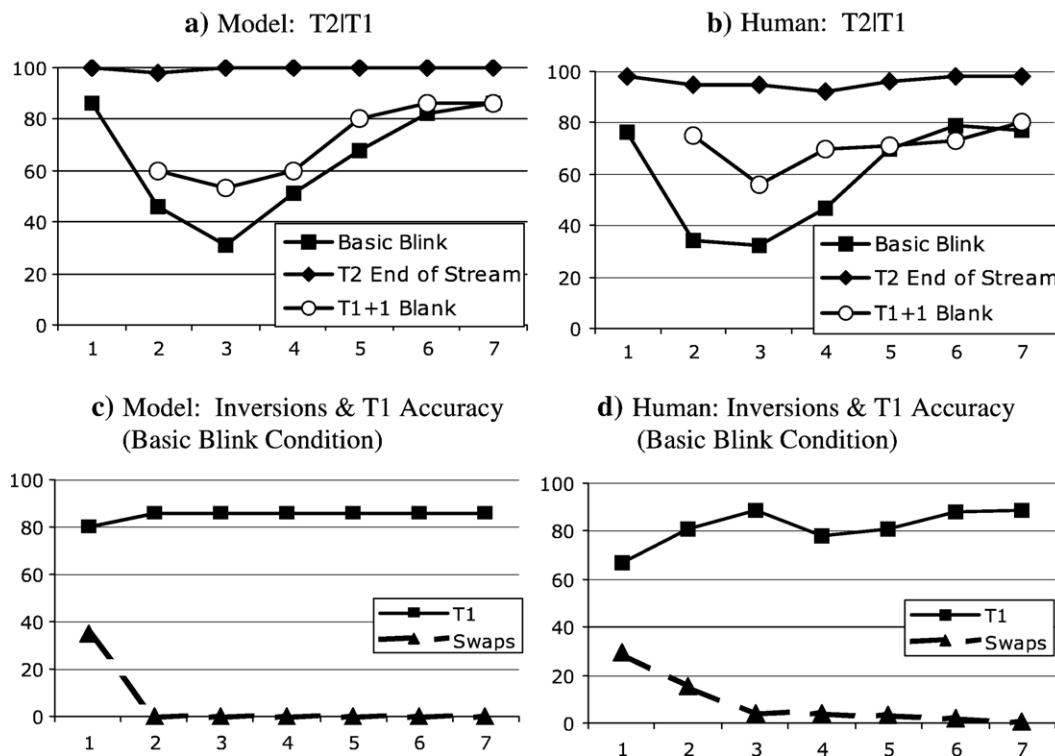


Fig. 3 – The  $ST^2$  model's performance (a, c) compared to human data (b, d). In all cases, a letters-in-digits task was considered with a 100 ms SOA. T2 performance (a, b) represents the accuracy in reporting T2 on trials in which T1 was reported. In panels c and d, the lines at the top of the graph show T1 accuracy, while the lines at the bottom denote the percent chance for the reported order of T1 and T2 to be inverted. Human data are from Chun and Potter (1995) except the T2 end of stream data, which is from Giesbrecht and Di Lollo (1998). Horizontal axes represent lag, while vertical axes denote accuracy. In the T1+1 blank condition there is no lag-1 case, since that slot is blank. Model data reproduced from Bowman and Wyble (2007). This diagram is reproduced from Bowman and Wyble (2007).

have not provided a data fit, the model is consistent with the finding of labile attention at short T1–T2 SOAs (Potter et al., 2002).

#### 4.1.3. Blink attenuation with T1+1 blank

The model demonstrates the key finding that inserting a blank in the T1+1 slot attenuates the blink; see Fig. 3(a). The sequence of events that generates this phenomenon is as follows. Firstly, an unmasked item yields a strong activation trace. Thus, a T1 followed by a blank generates a higher amplitude trace at the TFL. Secondly, and this is the critical step, tokenization, both in the sense of binding unit and token allocation, completes more quickly. This is because, through the previously discussed gate-trace mechanism, the model turns activation strength into encoding duration, on the principle that stronger evidence for an item (as encapsulated by activation strength) leads to more rapid encoding into WM. Thus, placing a blank after a T1 greatly shortens its tokenization. Thirdly, tokenization will be more likely to have finished before the T2 has decayed at the TFL, increasing the probability that the T2 will fire the blaster.

One subtlety of the T1+1 blank data is that performance is particularly elevated at lag-2. The ST<sup>2</sup> model obtains a qualitatively similar pattern at lag-2, since it is the only data point in which T2 is not strongly forward masked. (Although backward masking is far stronger, weak forward masking is also included in ST<sup>2</sup>.) Absence of forward masking increases the bottom-up trace strength of the T2, which gives it a small advantage according to the principles we discuss in the next section. However, it could be that lag-2 sparing in the T1+1 condition is actually a reflection of a form of spreading the sparing (as discussed in Section 2.2.9), which would explain the quantitatively smaller effect that the ST<sup>2</sup> model (which does not reflect spreading the sparing) currently obtains.

However, the key principle that blink attenuation with T1+1 blank reflects is that there is a reciprocal relationship between bottom-up trace strength and blink depth. This is obtained in ST<sup>2</sup> by tying tokenization time to trace strength.

#### 4.1.4. Blink attenuation with T2+1 blank

We compare the model to data from Giesbrecht and Di Lollo (1998), who examined the effect of placing T2 at the end of the stream; see Fig. 3(a). As required, the blink is obliterated. However, the ST<sup>2</sup> model obtains attenuation of the blink with T2 unmasking in a different manner to attenuation with T1 unmasking. Specifically, a T2+1 blank produces strong T2 traces that are more likely to outlive the blink. That is, T2 unmasking does not affect how long the blaster is held offline by T1 tokenization, but it does make the T2 more “resilient” to this blaster unavailability.


#### 4.1.5. Delayed T2 consolidation

Following on from the previous subsection, T2s encoded by ST<sup>2</sup> during the blink typically possess strong activation traces. (Variation in activation strength could reflect spontaneous noise or inherent differences in how particular targets stand out; Wyble and Bowman, 2005.) Consequently, T2s are often seen during the blink because they outlive blaster unavailability. As a result, in the average, T2s are consoli-


dated later during the blink. Fig. 4(a) shows new simulation results from the ST<sup>2</sup> model that encapsulates this effect, by measuring the peak of TFL activity for the T2 node. The T2 consolidation delay qualitatively mirrors the blink curve: T2 consolidation is most delayed when the blink is deepest and undelayed (i.e. at baseline) post blink recovery. Interestingly, the model suggests that T2 consolidation at lag-1 is accelerated, relative to post recovery baseline, which is consistent with the, previously discussed, increased processing of the T1+1 slot.

Although a human study which explores the full spectrum of lags is not available, there is solid evidence of delayed T2-

#### (a) ST<sup>2</sup> Model



#### (b) LC-NE Model



**Fig. 4 – T2 consolidation delay by lag for the ST<sup>2</sup> and LC-NE models.** For ST<sup>2</sup>, consolidation delay was measured as the time from target onset to peak amplitude of the TFL unit for that target. For LC-NE, consolidation delay was measured as the time from target onset to peak amplitude of the detection layer unit for that target. At each lag we averaged across all seen T2 trials. All ST<sup>2</sup> parameter settings were as presented in Bowman and Wyble (2007). Note though that even though the qualitative match between these two sets of data is very good, quantitatively they are very different, as reflected by them being presented on very different scales.

P3s during the blink<sup>6</sup> (Martens et al., 2006; Sessa et al., 2006; Vogel and Luck, 2002); see Section 2.2.5. This provides important support for the ST<sup>2</sup> model. In addition, confirmation of accelerated T2 consolidation at lag-1 would be particularly compelling evidence for the model; although, this is difficult to demonstrate since it is hard to isolate a T2 P3 at lag-1.

#### 4.1.6. Spatial specificity of lag-1 enhancement

Although the model has not yet been used to simulate spatial phenomena, from a theoretical perspective, the ST<sup>2</sup> blaster is spatially specific. That is, it is rapidly directed towards the location at which a salient item occurs. This is consistent with the heritage that underlies this mechanism, viz, transient spatial attention, as identified by Nakayama and Mackeben (1989). As previously discussed, this spatial specificity is consistent with findings in RSVP presentation settings (Visser et al., 1999; Wyble et al., submitted for publication-b).

#### 4.1.7. T1–T2 costs at lag-1

ST<sup>2</sup> does exhibit T1–T2 costs at lag-1. The loss in T1 accuracy at lag-1 (see Fig. 3(c)) arises since, when T2 is very strong and T1 weak, binding can complete before T1 is strongly active, yielding a successful binding from token 1 to T2 and a failed binding to T1. Moreover, at lag-1, the model is inaccurate at determining the order of the two targets; see Fig. 3(c). At lag-1, often, both T1 and T2 are bound to the first token. This reflects a loss of episodic information: due to the temporal proximity of T1 and T2, the model fails to allocate discrete episodic contexts for the two targets. These findings of costs at lag-1 are supported by human data; see Section 2.2.7.

#### 4.1.8. Late stage bottleneck

It is argued in Bowman and Wyble (2007, section 4.8) that ST<sup>2</sup> is consistent with a late bottleneck. In particular, none of the type layers in ST<sup>2</sup> is strongly suppressed during the blink. That is, types are always, at least to some extent, extracted from a target, regardless of whether it is blinked. However, typically, a T2 that is seen during the blink has a stronger trace than a missed T2, either by variation of input strength, or by being excited by the blaster.

#### 4.1.9. Spreading the sparing

In respect of the parameter settings used in Bowman and Wyble (2007), the ST<sup>2</sup> model does not replicate spreading the sparing. The inhibitory projection from the binding pool to the blaster (marked (c) in Fig. 1) is sufficiently strong that ongoing tokenization renders the blaster unavailable soon after T1 starts to be encoded. Thus, the model exhibits lag-1 sparing, but, in the context of a continuous sequence of targets and a standard SOA of around 100 ms, it does not exhibit lag-2 or -3 sparing. This is because, whether interleaved with distractors

or not, any target appearing in the 200 to 500 ms interval post T1 onset, will not be able to fire the blaster immediately.

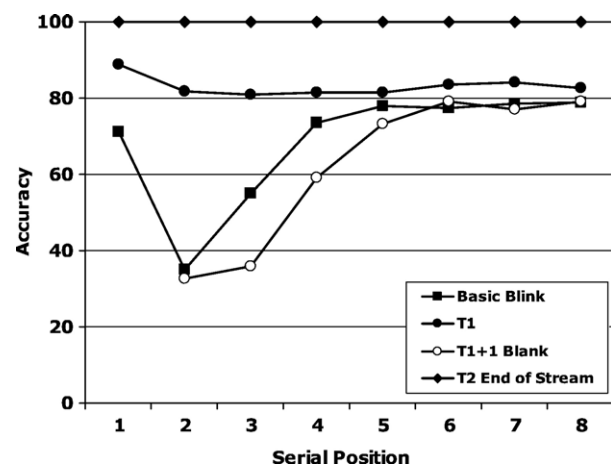
A revision of the model in which suppression of the blaster is not absolute and a somewhat more sophisticated token system is used, is under submission (Wyble et al., submitted for publication-a). In this revised model, a continuous stream of target-related activity can repeatedly re-fire the blaster, but at the expense of loss of episodic distinctiveness, e.g. order and repetition blindness errors (Wyble et al., submitted for publication-a). This revised model also replicates the finding that the blink can be reversed, in the sense that, although a T2 in the sequence T1 D T2 T3 would be blinked, the T3 would not be (Olivers et al., 2005). This is because the T2 overcomes blaster suppression; however, because of the difficulty of re-firing the blaster, T2 misses this benefit, which falls on the T3.

## 4.2. The LC–NE model

We now move to an assessment of the LC–NE model. The simulation results that we present are based upon the equations and parameter settings presented in Nieuwenhuis et al. (2005b), which we have re-implemented. Evidence that we have faithfully replicated the Nieuwenhuis et al. model is given by the fact that we generate a blink curve consistent with that found in Nieuwenhuis et al. (2005b). (Although, to maintain consistency with the presentation of ST<sup>2</sup> data, and indeed the human behavioral data, we use the conditional measure: T2 | T1. This though has little major effect on the shape of the basic blink curve.)

### 4.2.1. The basic blink

The LC–NE model generates a blink curve with lag-1 sparing; see Fig. 5. As previously suggested, the blink obtained follows the profile of the LC refractory period and lag-1 is spared, since it benefits from the NE release arising from the T1-induced LC firing.



**Fig. 5 – Standard blink conditions generated from the re-implemented LC–NE model. Accuracy is T2 | T1 for all conditions except T1. For T1 + 1 blank, lag-1 is not a valid data point, since it contains the blank. For full details of conditions see main text.**

<sup>6</sup> Although the quantitative match between ST<sup>2</sup>'s T2 consolidation latencies and human P3 latencies is surprisingly close, the effect size (i.e., difference between consolidation latencies inside the blink and post recovery) suggested by the model is somewhat larger than in the human data (Martens et al., 2006). However, a full investigation of this issue awaits further ERP studies and a targeted effort to generate P3s from the model.

#### 4.2.2. Increased processing of T1+1 slot

The T1-induced LC firing benefits the T1+1 slot. Furthermore, the enhancement is item non-specific, in the sense that it would also benefit a distractor in the T1+1 slot. Thus, the model should be viewed as consistent with the finding of increased priming from distractors following T1s (Chua et al., 2001). Nieuwenhuis et al.'s model is also consistent with a temporal, rather than sequential (by-item) interpretation of blink onset. That is, it suggests that the T1-induced enhancement has a minimal extent, which is broadly fixed and is not regulated by intervening distractors. This is supported by a number of findings; see Section 2.2.2.

#### 4.2.3. Blink attenuation with T1+1 blank

In the LC-NE model (Nieuwenhuis et al., 2005b), the length of the LC refractory period is not fixed. Rather, stronger LC firings yield a longer refractory period. This raises a problem, since strength of LC firing is determined by the strength of the target that drives it. Thus, greater bottom-up trace strength leads to a longer refractory period. This facet of Nieuwenhuis et al.'s model has the consequence that unmasking T1s (i.e. T1+1 blank conditions) deepens and lengthens the blink.

We have explored this issue with simulation runs. As acknowledged by the authors, the LC-NE simulation presented in Nieuwenhuis et al. (2005b) does not model masking. Thus, we have had to indirectly model the effects of following a T1 by a 100 ms blank instead of a distractor. Specifically, we model an RSVP sequence containing a T1 followed by a blank as .... D T1 T1 D ....., where a standard (nonblanked) T1 is modelled as .... D T1 D D .... (Ds denote arbitrary distractors). (See Appendix A, Section A.2, for discussion of the technical issues associated with this manipulation.) This approach is consistent with the observation that the after-image of a visual stimulus remains for some hundreds of milliseconds if a masking item does not follow (see, e.g. the iconic memory literature (Coltheart, 1983; Sperling, 1960) or single cell recordings of persistence in the visual pathway (Keysers and Perrett, 2002; Keysers et al., 2005)).

The results of our simulations are shown in Fig. 5. As predicted, the model shows a deeper and longer blink for the T1+1 blank condition and this is indeed because stronger LC firing (arising from stronger T1s) leads to a longer refractory period. This weakness of the model was acknowledged in Nieuwenhuis et al. (2005b).

#### 4.2.4. Blink attenuation with T2+1 blank

The LC-NE model does generate blink attenuation with a T2+1 blank. In order to match T2 end of stream data, we investigated a similar manipulation to that just discussed, but now for T2. Specifically, we compared sequences of the form .... D T1 D D T2 D .... (basic blink, lag-3) with those of the form .... D T1 D D T2 T2 T2 .... (T2 End of Stream, lag-3). As found with humans (see Fig. 3(b)), the model exhibited ceiling performance, see Fig. 5. As was the case with the ST<sup>2</sup> model, the blink is removed with T2 unmasking, since high amplitude T2s benefit from higher intrinsic strength and a greater ability to outlive the LC refractory period.

#### 4.2.5. Delayed T2 consolidation

T2 consolidation latencies are indeed delayed during the blink. This can be seen in Fig. 4(b),<sup>7</sup> which shows the results of generating T2 consolidation latencies from our reimplemented LC-NE model. Furthermore, as was the case for ST<sup>2</sup>, the LC-NE model generates accelerated P3 consolidation latencies (relative to recovery baseline) at lag-1. This is again because the lag-1 item benefits from the T1-induced enhancement, which here amounts to NE release generated by the T1-induced LC firing. In fact, the LC-NE model generates the same qualitative pattern of T2 consolidation latencies by lag as the ST<sup>2</sup> model, which gives strong credence to this particular theoretical prediction.

#### 4.2.6. Spatial specificity of lag-1 enhancement

The LC-NE enhancement is assumed to be completely generalized, both featurally and spatially. That is, any item, whatever its features or spatial location, would be enhanced. In other words, the LC-NE provides a purely temporal filter. As acknowledged by Nieuwenhuis et al., this means that the LC-NE model cannot explain the spatial specificity of lag-1 sparing, at least, without assuming a further mechanism to provide that specificity (Nieuwenhuis et al., 2005b); see Section 2.2.6.

#### 4.2.7. T1-T2 costs at lag-1

The LC-NE model does not generate T1-T2 costs at lag-1. Firstly, as can be seen in Fig. 5, there is no decrement in T1 performance at lag-1.<sup>8</sup> In addition, the model does not encode order information. Thus, the finding of increased order inversions at lag-1 cannot be investigated.

#### 4.2.8. Late stage bottleneck

It is difficult to be definitive with regard to the stage of the bottleneck in the LC-NE model, since one would need to postulate a more detailed relationship between layers of the model and stages of processing. Of most relevance would be the attribution of semantic processing to a particular layer in the model. In such an abstract model as the LC-NE model, such attribution is not obvious. This said, a comparison of the activation traces generated by seen and missed T2s does not obviously suggest an early bottleneck. In particular, seen and missed T2s generate almost identical decision layer traces, while detection layer traces are more different, although, even missed T2s do elicit nontrivial detection layer activation. Furthermore, the theoretical underpinnings of

<sup>7</sup> If this pattern of data were to be related to P3 latencies, then the quantitative match would not be perfect. In particular, latencies are much shorter here than that suggested by P3s. In addition, the effect size (i.e., difference between consolidation delays inside and outside the blink) is smaller than that suggested by P3 data (Martens et al., 2006).

<sup>8</sup> In fact, as should be evident from Fig. 5, the model has a tendency to produce enhanced T1 performance at lag-1. This is due to a technical detail concerning the handling of distractors, which means that targets have lower decision layer activations than distractors. As a result, a T1 followed by a target (the T2) receives less suppression at the decision layer than a T1 followed by a distractor. This though is an implementation detail, which we consider separate from the theoretical position of the LC-NE model.

the LC–NE hypothesis are also consistent with a late bottleneck. For example, [Shea-Brown et al. \(submitted for publication\)](#) suggests that a primary role of the LC–NE system is to facilitate post decision processing. Thus, as far as it is possible to say, the LC–NE model is consistent with a late stage bottleneck.

#### 4.2.9. Spreading the sparing

The LC–NE model does not generate spreading the sparing. This is evident in the lag-2 data point of the T1+1 blank condition, see [Fig. 5](#). As previously discussed, this data point is modelled as a stream of distractors containing the subsequent T1 T1 T2, i.e. a continuous sequence of targets. For this data point, the model performs below the basic blink curve (see [Fig. 5](#)), while humans are almost at baseline performance (see [Fig. 3\(b\)](#)).

Because of the nature of the LC refractory period, the LC–NE model fails to spread the sparing. This is because, firstly, if a T1 fires the LC, it will go into a refractory period (in this sense it is ballistic) and, secondly, ongoing bottom-up activation to the LC (as generated by a continuous stream of targets) cannot overcome the refractory period. More specifically, our investigations with the LC–NE model suggest that when presented with a continuous stream of particularly strong targets a robust refractory period remains (even though its shape may change somewhat).

#### 4.3. Discussion

Beyond the realm of specific behavioral phenomena, there are a number of other important differences between the ST<sup>2</sup> and LC–NE approaches. Firstly, ST<sup>2</sup>'s TAE is an additive enhancement, while the LC enhances by increasing the gain of the activation function. Thus, Nieuwenhuis et al.'s mechanism is multiplicative, which yields a gating aspect that is not present with ST<sup>2</sup>'s TAE/blaster. ST<sup>2</sup>'s additive enhancement has the virtue of simplicity; in particular, it does not require any mechanisms that are not present in the most standard neural networks framework. However, although the additive bias approach works well in the ST<sup>2</sup> context, in which the enhancement is very brief, we have had trouble modeling spreading the sparing with this mechanism. Spreading the sparing suggests a more sustained amplification or, at least, a rapid series of attention pulses. In this context, an additive bias is susceptible to pushing even stimuli never presented over threshold, due to the accumulation over time of additive excitation. As a result, our revision of the ST<sup>2</sup> framework to handle spreading the sparing has moved to a multiplicative gain ([Wyble et al., submitted for publication-a](#)). Thus, in this respect, the two models are converging.

Secondly, as acknowledged by the authors, Nieuwenhuis et al.'s approach does not sustain a memory trace to the end of a trial. That is, T1 and T2 activations rise and fall at the detection layer with a time-course in the range of a few hundred milliseconds of simulated time. WM maintenance is viewed as a separate mechanism that is beyond the scope of the LC–NE model of [Nieuwenhuis et al. \(2005b\)](#).

This section has demonstrated that both the ST<sup>2</sup> and LC–NE models generate a number of the key empirical phenomena. However, both (as currently formulated) have difficulty with

spreading the sparing and the LC–NE model additionally has difficulties replicating blink attenuation with T1+1 blank, T1–T2 costs at lag-1 and requires the assumption of an additional mechanism to explain the spatial specificity of the lag-1 enhancement. Since we believe it is such a central data point, in the next section we discuss how the LC–NE model could be extended in order to model the reciprocal relationship between bottom-up trace strength and the AB bottleneck. This will have the additional benefit of adding a WM maintenance mechanism to the model. Furthermore, we discuss what implication such additions would suggest for LC neurophysiology.

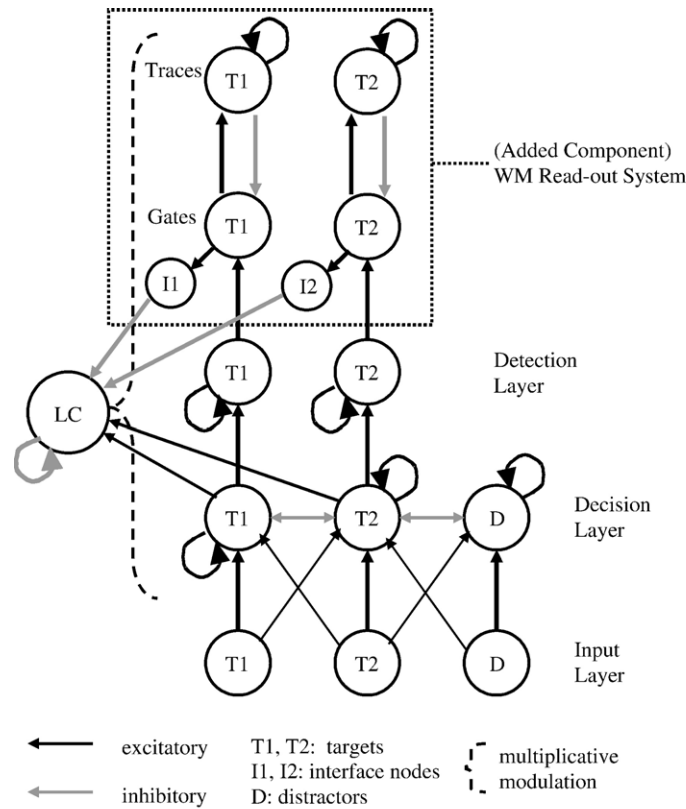
### 5. Extending the LC–NE model with gate–trace pairs

Our intent is not to develop a fully formed alternative blink model. Rather, this section simply considers a possible extension of the LC–NE model that would enable it to simulate blink attenuation with T1+1 blank and maintain items in WM beyond encoding. The mechanisms that we consider are inspired by those already present in ST<sup>2</sup>. This is undertaken as an exploratory exercise that, in particular, could inform further neurophysiological investigations of the LC–NE system, especially of the major projections between the cortex and LC. It thus also clarifies the neurophysiological implications of the cognitive-level mechanisms proposed by ST<sup>2</sup>.

In the LC–NE network, transient representations of decision results arise at the detection layer, which is where the LC–NE implementation of [Nieuwenhuis et al. \(2005b\)](#) finishes. However, one could also imagine a WM encoding system that “reads out” from these detection layer activations. Such a system could be based upon the gate–trace mechanisms we have highlighted in the ST<sup>2</sup> context (see Section 3.1.3). We explore this possibility in the context of the LC–NE model. However, it is important to note that in this exploration we leave the basic LC–NE model unchanged; thus, we highlight a strict extension of the LC–NE system of [Nieuwenhuis et al. \(2005b\)](#). The structure of the extended system is shown in [Fig. 6](#) and implementation details can be found in [Appendix A](#).

As is inherent in the gate–trace approach, the trace nodes maintain items in WM, while gate nodes enable items to be encoded into WM, i.e. they gate access to trace nodes. The gate–trace extension behaves as follows. Firstly, activation onset of gate nodes mirrors detection node activations, subject to a small time delay. Secondly, active gate nodes drive their corresponding trace node until it crosses threshold, at which point encoding is deemed to have completed. As a result, the trace node rapidly suppresses its gate and enters a self-sustaining attractor state at which point the target has been successfully encoded into WM.

Thus, the gate–trace extension adds the capacity to hold items in WM once they have been successfully encoded. However, in addition, we have investigated tying LC suppression to ongoing encoding, according to the ST<sup>2</sup> principle that withholding of the attentional enhancement should be coupled to length of WM encoding. Thus, we have also added a projection from gate nodes to the LC. Overall, this has an inhibitory effect on the LC, although it passes through



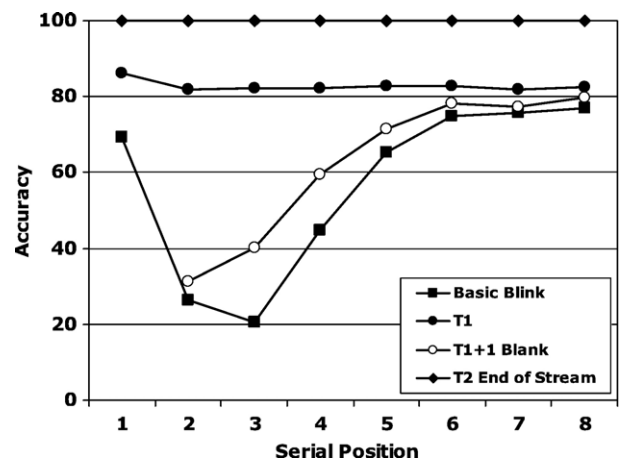
**Fig. 6 – Extended LC-NE model. Note, inhibitory and crosstalk connections between T1 and D are not shown for simplicity of presentation. Interface nodes, interface the gate–trace system with the LC, ensuring gradual interactions between the components.**

an intermediate node on the way. We call these intermediate nodes interface nodes, since they interface gate nodes with the LC. They are needed to turn the sharp changes in activation that occur at the gates into more gradual effects on the LC. Sharp discontinuities of input to the LC, whether they be excitatory or inhibitory, disrupt the sensitive balance between LC state and NE level, generating spurious changes in LC state. This is particularly the case with regard to the offset of activation of the T1 gate, which (due to trace node suppression) is very rapid.

The behavior of this system with regard to the key phenomena is shown in Fig. 7. For these data, a target was considered as “seen” if its corresponding trace node was in its attractor state at the end of the simulated trial. (In contrast, the detection unit crossing threshold counted the target as “seen” in the original LC-NE model; Nieuwenhuis et al., 2005b.)

Thus, the extended LC-NE model generates an interesting profile of data. Firstly, it generates a longer basic blink curve, which is more consistent with human data. Secondly, attenuation of the blink with T2+1 is preserved. Thirdly, as anticipated, the model obtains blink attenuation with T1+1 blank. The lag-2 data point remains a difficulty with the LC-NE model, as performance can only recover, at best, to the level defined by the refractory period, no matter how rapidly the T1 is encoded. This difficulty would be partially offset by an implementation of weak forward masking, as was done in the ST<sup>2</sup> model (Bowman and Wyble, 2007). Furthermore, as suggested pre-

viously, the good performance obtained at lag-2 in the T1+1 blank condition, could be attributed to a spreading the spring-like effect since, as we argue, the T1+1 blank is akin to sustaining an iconic representation of T1, yielding a continuous sequence of target-related activity.



**Fig. 7 – Standard blink conditions generated from the extended LC-NE model. Accuracy is T2 | T1 for all conditions except T1. For T1+1 blank, lag-1 is not a valid data point, since it contains the blank. For full details of conditions see main text.**

The extended LC–NE model does generate blink attenuation with T1+1 blank. However, a number of phenomena remain problematic, e.g. T1–T2 costs at lag-1 and spreading the sparing. A full consideration of the former of these would require, first, a careful consideration of T1–T2 competition at lag-1 and, second, the addition of something like the ST<sup>2</sup> binding pool and token system, in order to investigate order errors (which as a by-product would introduce a means to represent repetitions). Such investigations are beyond the scope of this article.

Spreading the sparing is a difficulty for all published models (although, a revised ST<sup>2</sup> model that handles the effect is under submission (Wyble et al., submitted for publication-a)). However, this phenomenon is particularly tricky for the LC–NE framework to handle; we return to this issue in the final discussion, as we do to the question of the spatial specificity of the lag-1 enhancement.

However, with regard to the central theme of this article, this section has shown how the LC–NE framework could be extended to model blink attenuation with T1+1 blank by incorporating an ST<sup>2</sup>-inspired gate–trace mechanism. In fact, such a mechanism would be consistent with the more general finding of modulation of blink depth by ease of the T1 task (Jolicoeur, 1998; Wyble and Bowman, 2005).

Finally, what then would be the implications of the extensions introduced here for neurophysiology? Firstly, there is nothing controversial with regard to gate–trace circuits, since inhibitory inter-neurons are common-place in the brain (O'Reilly and Munakata, 2000). Thus, the central point of uncertainty is the suggested projection from WM encoding areas to the LC, which should have an inhibitory effect. Although it is known that there are major projections from frontal areas (especially orbital frontal cortex and the anterior cingulate) to the LC (Aston-Jones and Cohen, 2005), whether these have the required characteristics to support the extended LC system remains to be answered. For an approach such as that suggested in this section to obtain greater credence, projections such as these would need to be identified in the primate brain.

## 6. Final discussion

The ST<sup>2</sup> model provides a concrete and broad scope theory of the AB, which matches a large spectrum of empirical phenomena. In particular, the model generates a blink curve with lag-1 sparing, increased processing of the T1+1 slot, blink attenuation with T1+1 blank, blink attenuation with T2+1 blank, delayed consolidation for T2s seen during the blink and T1–T2 costs at lag-1. In addition, the mechanisms postulated are consistent with the known spatial specificity of the lag-1 enhancement and a late stage blink bottleneck. The major outstanding phenomenon that the currently published implementation of the model does not address is spreading the sparing.

Nieuwenhuis et al's LC–NE model suggests a compelling theory for the cause of the AB based on a neurophysiologically prescribed theory of attentional function (Aston-Jones and Cohen, 2005; Gilzenrat et al., 2002; Nieuwenhuis et al., 2005a). The model also generates a number of the key AB

phenomena; for example, it generates a blink curve with lag-1 sparing, increased processing of the T1+1 slot, blink attenuation with T2+1 blank and delayed consolidation for T2s seen during the blink. It is also broadly consistent with a late stage bottleneck. However, it does not generate blink attenuation with T1+1 blank, T1–T2 costs at lag-1 or spreading the sparing and (without the assumption of a further mechanism) the LC theory does not generate a spatially specific lag-1 enhancement.

While the ST<sup>2</sup> model is neurophysiologically plausible, in the sense that it is formulated using known neurobiological mechanisms, there is uncertainty concerning the exact localisation of some of ST<sup>2</sup>'s components (although, see Bowman and Wyble (2007, section 4.9) for a proposal in this respect). Thus, an exploration of the relationship between Nieuwenhuis et al's brain-level proposal and ST<sup>2</sup>'s cognitive-level proposal is valuable. This is what we have considered in this article.

It should also, though, be acknowledged when comparing the two models that they have somewhat different intent and scope. ST<sup>2</sup> is a more elaborate model than the LC–NE model, containing more layers and components. This reflects ST<sup>2</sup>'s intent to be a relatively broad scope model of temporal attention and WM encoding and maintenance. In contrast, the LC–NE model, as presented in Nieuwenhuis et al. (2005b), does not claim to model WM; rather its value lies with the fact that a blink effect is obtained despite only adding a minimal set of additional assumptions to those included in previous LC–NE models (Gilzenrat et al., 2002).

### 6.1. The blaster and the LC

The LC enhancement and ST<sup>2</sup>'s TAE/blaster have a number of similarities, e.g. both are initiated by detection of a salient stimulus, they are type non-specific and their temporal profiles are similar. However, there are important differences between the two.

Firstly, the blaster is postulated to have its effect on stimulus representations relatively early in the processing pathway, certainly no later than inferotemporal cortex and perhaps actually in visual cortex (Bowman and Wyble, 2007). The hypothesis being that “blasted” types obtain greater bottom-up trace strength, giving them an encoding/tokenization advantage. Thus, in ST<sup>2</sup>, the attentional gate works by regulating bottom-up trace strength. In contrast, the LC is suggested to have its main effect on decision and response systems (Aston-Jones and Cohen, 2005). Furthermore, LC innervation is not dense in visual cortex (especially primary visual cortex) (Nieuwenhuis et al., 2005a) and LC innervation of the temporal lobes is more focused on the superior temporal gyrus (especially, the Temporo-parietal Junction (TPJ)) than the inferior temporal gyrus.

Secondly, the blaster is assumed to be spatially specific; however, the LC enhancement would be expected to be spatially general, as suggested by neuroanatomical studies of the pattern of noradrenergic projections (Nieuwenhuis et al., 2005a). The LC–NE approach cannot resolve this difficulty without introducing a further mechanism that enhances specific locations, and that mechanism is likely to

be similar to ST<sup>2</sup>'s blaster. One speculative (and perhaps less than parsimonious) explanation could be that the LC provides the temporal profile of a transient form of attention, such as produced by the blaster. One of the candidate areas for locating the blaster is the TPJ (Bowman and Wyble, 2007) and the LC is known to strongly innervate the TPJ. If one assumed a spatially specific amplifier at the TPJ, then the LC could be “amplifying the amplifier” in a transient fashion.

## 6.2. Correlates of the P3

An important question is how to relate models to the P3. The LC–NE theory suggests a specific neural substrate for the P3 (Nieuwenhuis et al., 2005a). The link from activation traces in the LC–NE model to the P3 though is not clear-cut. In particular, all the target-induced activations in the LC–NE model peak and indeed finish a good deal earlier than the known time-course of the P3, as is evident in Fig. 4(b). Thus, these target-induced activations are not direct analogues of the P3. This leaves the possibility that the P3 is an indirect (delayed) consequent of these target-induced activations; that is, that there is a latency offset between model activation traces and what is observed as the P3. However, assessment of this explanation awaits a concrete proposal for the mechanics of this latency offset. Although not perfect, the time-course of ST<sup>2</sup> target activation (especially at the TFL) is more consistent with that of the P3, which, in an RSVP setting peaks between 400 and 450 ms after the onset of the eliciting stimulus. Ongoing work in our lab is directly addressing the issue of how to generate P3 profiles (and other ERP components) from the ST<sup>2</sup> model.

## 6.3. Bottom-up trace strength and blink depth

This article has particularly attempted to emphasize the reciprocal relationship between bottom-up trace strength and blink depth and the importance this phenomenon has for models of the AB. In particular, this relationship falsifies many of the simpler AB models that one could imagine. Most of the obvious ways to obtain a blink-like effect would imply that the size of the impairment would be directly proportional to the size of the T1 bottom-up activation trace, i.e. stronger T1 activations would generate deeper blinks. For example, a very simple way to obtain a blink-like impairment would be via a layer of mutually inhibitory item nodes. Thus, while the T1 was active, it would suppress any T2 at that layer. Along with a number of other difficulties (c.f. Bowman and Wyble, 2007), this approach suggests that stronger T1s would more strongly inhibit T2s and thus, stronger T1s would generate deeper blinks. In contrast, in ST<sup>2</sup>, stronger T1s are encoded more rapidly, which means that T2s have to wait less before starting encoding. Any competitor neural model needs to postulate a means to obtain this key reciprocal relationship.

## 6.4. Spreading the sparing

Modeling spreading the sparing represents the most significant challenge in the AB literature. In the context of the ST<sup>2</sup>

and LC–NE models, at the least, this phenomenon suggests that unavailability of the attentional enhancement (blasting or NE release respectively) is not absolute and that an unbroken stream of salient stimuli should be able to overcome this unavailability. The phasic mode LC firing hypothesis inherent in the LC–NE model seems particularly difficult to reconcile with this requirement, as spreading the sparing stands against the existence of an (effectively) uncounteractable refractory period. The parameter settings in Nieuwenhuis et al. (2005b) ensure that the refractory period cannot be overcome through an unbroken stream of target-related activity.

In contrast to the notion that an inherent refractory period of attentional allocation is the cause of the blink, the ST<sup>2</sup> model suggests that the cause is a deliberate suppression of attention. The advantage of this account is demonstrated in the model of Wyble et al. (submitted for publication-a), which gives ongoing target activity the capacity to overcome suppression of the blaster. As a result, a continuous sequence of targets can be encoded “together” into WM. However, this increased identity encoding comes at the cost of a decline in episodic distinctiveness, which manifests as a loss of temporal order information and increased repetition blindness within a continuous sequence of targets.

## Acknowledgments

The ideas presented in this article have benefited significantly from discussions with Mary Potter, Phil Barnard, Su Li, Mark Nieuwenstein, Chris Olivers, Martijn Meeter, Sander Martens, Marius Usher, John Taylor, Jonathan Cohen and Eddy Daveelaar. We would also like to thank one anonymous reviewer for valuable insights.

## Appendix A. Implementation details

### A.1. The extended LC–NE model

Pairs of gate–trace nodes for each target stimulus, connected to their corresponding detection layer nodes, form the “read out” system for WM encoding. In addition, interface nodes translate gate node activation into inhibition of the LC–NE system. The state  $Z$  of a node is updated once every iteration, by numerically integrating the ordinary differential equation for that node using a simple Euler method and then computing the activity of the node using its sigmoidal activation function  $f(Z)$  which is defined as,

$$f(Z) = \frac{1}{1 + e^{-g(Z-b)}}$$

where  $g$  is the multiplicative gain and  $b$  is the bias. The dynamic effect of the noradrenergic output  $u$  on  $g$  is modelled as

$$g = B + ku$$

where  $B$  is the base level of gain and  $k$  is a scaling constant fixed at 1.5.



The differential equation for computing the state of a gate node  $G$  is given by

$$\frac{dX_G}{dt} = -X_G + w_{GT}[f(X_T) - \theta]_+ + w_{GR}f(X_R)$$

where  $X_G$  is the state of  $G$ ,  $w_{GT}$  (4) is the weight of the link from detection layer node  $T$  to  $G$ ,  $w_{GR}$  (-10) is the weight of the link from  $G$ 's trace node  $R$  to  $G$ ,  $\theta$  (0.532) is a threshold set to the baseline activity level of  $T$  and  $[x]_+ = \max(x, 0)$ . The equation above is used to compute the state of  $G$  only after  $f(X_T)$  crosses the threshold of 0.67 used in the original LC-NE model. The activation function for  $X_G$ , i.e.  $f(X_G)$ , uses a base level gain ( $B$ ) of 1 and a bias ( $b$ ) of 1.

The equation for computing the state of a trace node  $R$  is given by:

$$\frac{dX_R}{dt} = -X_R + w_{RG}f(X_G) + w_{RR}f(X_R)$$

where  $X_R$  is the state of  $R$ ,  $w_{RG}$  (0.2) is the weight of the connection from the gate node  $G$  to  $R$  and  $w_{RR}$  (3) is the weight of the self-loop on  $R$ . The activation function for  $X_R$ , i.e.  $f(X_R)$ , uses a base level gain ( $B$ ) of 20 and a bias ( $b$ ) of 0.3.

The equation for computing the state of an interface node  $I$  is given by:

$$\tau_I \frac{dX_I}{dt} = -X_I + w_{IG} \min(f(X_G), \mu)$$

where  $X_I$  is the state of  $I$ ,  $w_{IG}$  (20) is the weight of the connection from the gate node  $G$  to  $I$ ,  $\tau_I$  (10) is the time constant of  $X_I$  and  $\mu$  (0.8) is an upper threshold on the activity level of  $G$  that is received by  $I$ . The activation function for  $f(X_I)$ , uses a base level gain ( $B$ ) of 4 and a bias ( $b$ ) of 1.4.

In the extended model, the LC-NE system receives inhibitory input from the interface nodes. Consequently, the equation modeling the LC state variable  $v$  is modified to:

$$\tau_v \frac{dv}{dt} = w_{vX}[f(X_{T1}) + f(X_{T2})] + w_{vI}([f(X_{I1}) - \rho]_+ + [f(X_{I2}) - \rho]_+) + v(a - v)(v - 1) - u$$

where  $w_{vI}$  (-0.4) represents the weights of the inhibitory connections from the interface nodes  $I1$  and  $I2$  with activities  $f(X_{I1})$  and  $f(X_{I2})$ , respectively, and  $\rho$  (0.8) represents a lower threshold on the activity of the interface nodes.  $w_{vX}$  (0.3) is the link weight from decision layer to LC and  $X_{T1}$  and  $X_{T2}$  are the states of the decision units.

All other equations governing the LC-NE system, the parameters therein and the connection weights as configured in the behavioral network of the original LC-NE model remain unchanged. Only the number of iterations comprising a simulation trial is increased from 2200 to 2400, to extend the trial over 48 units of model time and incorporate 8 lag positions for  $T2$ .

## A.2. Simulating blanks in the LC-NE model

Consideration of the  $T1+1$  blank condition is confused by the fact that placing an actual blank after the  $T1$ , i.e. presenting ....  $D T1\_D$  .... to the model, obliterates the blink. However, this is not due to absence of backward masking and indeed

would be inconsistent with the afterimage persistence of visual stimuli (Keyser and Perrett, 2002; Keyser et al., 2005). The blink is attenuated with ....  $D T1\_D$  .... because the  $T2$  benefits from reduced inhibitory competition. The blank interval in the  $T1+1$  position allows the activation level of the distractor node to fall to nearly resting level. In the LC-NE model, a  $T$  in a sequence ....  $D\_T D D$  .... (or, indeed, ....  $D\_D T D$  ....., etc.) will always have an advantage over a  $T$  in a sequence not containing a blank (e.g. ....  $D D T D D$  ....). Thus, blanks placed anywhere in the stream generate a forward going (in the sense of the stream) competitive advantage. However, there are a number of reasons why this cannot serve as an explanation of AB attenuation with  $T1+1$  blank. For example, if the LC-NE decision layer set-up was taken as a model of target — blank effects in the AB, it would predict that  $T1-1$  blank should massively improve  $T1$  performance and attenuate the blink, which it does not (Breitmeyer et al., 1999).

## REFERENCES

- Anderson, A.K., 2005. Affective influences on the attentional dynamics supporting awareness. *J. Exp. Psychol. Gen.* 134 (2), 258–281.
- Anderson, A.K., Phelps, E.A., 2001. Lesions of the human amygdala impair enhanced perception of emotionally salient events. *Nature* 411 (6835), 305–309.
- Aston-Jones, G., Cohen, J.D., 2005. An integrative theory of locus coeruleus–norepinephrine function: adaptive gain and optimal performance. *Annu. Rev. Neurosci.* 28, 403–450.
- Aston-Jones, G., Rajkowski, J., Cohen, J., 2000. Locus coeruleus and regulation of behavioral flexibility and attention. *Prog. Brain Res.* 126, 165–182.
- Battye, G., 2003. Connectionist Modelling of Attention and Anxiety. Unpublished PhD, The Medical Research Council's Cognition and Brain Sciences unit, Cambridge.
- Bowman, H., Wyble, B., 2007. The simultaneous type, serial token model of temporal attention and working memory. *Psychol. Rev.* 114 (1), 38–70.
- Braver, T.S., Cohen, J.D., 2000. On the control of control: the role of dopamine in regulating prefrontal function and working memory. In: Monsell, S., Driver, J. (Eds.), *Attention and Performance XVIII: Control of Cognitive Processes*. MIT Press, pp. 713–737.
- Breitmeyer, B.G., Ehrenstein, A., Pritchard, K., Hiscock, M., Crisan, J., 1999. The roles of location specificity and masking mechanisms in the attentional blink. *Percept. Psychophys.* 61 (5), 798–809.
- Chartier, S., Cousineau, D., Charbonneau, D. (2004). A Connectionist Model of the Attentional Blink Effect During a Rapid Serial Visual Task. In *ICCM 2004, International Conference on Cognitive Modelling*.
- Chua, F.K., Goh, J., Hon, N., 2001. Nature of codes extracted during the attentional blink. *J. Exp. Psychol.: Hum. Percept. Perform.* 27 (5), 1229–1242.
- Chun, M.M., 1997a. Temporal binding errors are redistributed by the attentional blink. *Percept. Psychophys.* 59 (8), 1191–1199.
- Chun, M.M., 1997b. Types and tokens in visual processing: a double dissociation between the attentional blink and repetition blindness. *J. Exp. Psychol. Hum. Percept. Perform.* 23 (3), 738–755.
- Chun, M.M., Potter, M.C., 1995. A two-stage model for multiple target detection in rapid serial visual presentation. *J. Exp. Psychol. Hum. Percept. Perform.* 21 (1), 109–127.

- Chun, M.M., Potter, M.C., 2000. The attentional blink and task switching within and across modalities. In: Shapiro, K. (Ed.), *Temporal Constraints in Human Information Processing*. Oxford Univ. Press, Oxford.
- Coltheart, M., 1983. Iconic memory. *Philos. Trans. R. Soc. Lond., B Biol. Sci.* 302 (1110), 283–294.
- Cowan, N., 2001. The magical number 4 in short-term memory: a reconsideration of mental storage capacity. *Behav. Brain Sci.* 24 (1), 87–114.
- Davelaar, E.J., Usher, M., 2004. An extended buffer model for active maintenance and selective updating. In: Bowman, H., Labiouse, C. (Eds.), *Proceedings of the 8th Neural Computation and Psychology Workshop*, vol. 15. World Scientific, Singapore, pp. 3–14.
- Dehaene, S., Sergent, C., Changeux, J.P., 2003. A neuronal network model linking subjective reports and objective physiological data during conscious perception. *Proc. Natl. Acad. Sci. U. S. A.* 100 (14), 8520–8525.
- Dehaene, S., Changeux, J.-P., Naccache, L., Sackur, J., Sergent, C., 2006. Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends Cogn. Sci.* 10 (5), 204–211.
- Desimone, R., Duncan, J., 1995. Neural mechanisms of selective visual attention. *Annu. Rev. Neurosci.* 18, 193–222.
- Di Lollo, V., Kawahara, J., Shahab-Ghorashi, S.M., Enns, J.T., 2005. The attentional blink: resource depletion or temporary loss of control? *Psychol. Res.* 69 (3), 191–200.
- Enns, J.T., Di Lollo, V., 2000. What's new in visual masking? *Trends Cogn. Sci.* 4 (9), 345–352.
- Fragopanagos, N., Kockelkoren, S., Taylor, J.G., 2005. A neurodynamic model of the attentional blink. *Cogn. Brain Res.* 24 (3), 568–586.
- Frank, M.J., Loughry, B., O'Reilly, R.C., 2001. Interactions between frontal cortex and basal ganglia in working memory: a computational model. *Cogn. Affect. Behav. Neurosci.* 1 (2), 137–160.
- German, D.C., Walker, B.S., Manaye, K., Smith, W.K., Woodward, D.J., North, A.J., 1988. The human locus coeruleus: computer reconstruction of cellular distribution. *J. Neurosci.* 8 (5), 1776–1788.
- Giesbrecht, B., Di Lollo, V., 1998. Beyond the attentional blink: visual masking by object substitution. *J. Exp. Psychol. Hum. Percept. Perform.* 24 (5), 1454–1466.
- Gilzenrat, M.S., Holmes, B.D., Rajkowski, J., Aston-Jones, G., Cohen, J.D., 2002. Simplified dynamics in a model of noradrenergic modulation of cognitive performance. *Neural Netw.* 15 (4–6), 647–663.
- Jolicoeur, P., 1998. Modulation of the attentional blink by on-line response selection: evidence from speeded and unspeeded task1 decisions. *Mem. Cogn.* 26 (5), 1014–1032.
- Kanwisher, N.G., 1987. Repetition blindness: type recognition without token individuation. *Cognition* 27 (2), 117–143.
- Kanwisher, N.G., 1991. Repetition blindness and illusory conjunctions: errors in binding visual types with visual tokens. *J. Exp. Psychol. Hum. Percept. Perform.* 17 (2), 404–421.
- Keyers, C., Perrett, D.I., 2002. Visual masking and RSVP reveal neural competition. *Trends Cogn. Sci.* 6 (3), 120–125.
- Keyers, C., Xiao, D.-X., Foldiak, P., Perrett, D.I., 2005. Out of sight but not out of mind: the neurophysiology of iconic memory in the superior temporal sulcus. *Cogn. Neuropsychol.* 22 (3/4), 316–332.
- Kranczoch, C., Debener, S., Engel, A.K., 2003. Event-related potential correlates of the attentional blink phenomenon. *Brain Res. Cogn. Brain Res.* 17 (1), 177–187.
- Luck, S.J., Vogel, E.K., Shapiro, K.L., 1996. Word meanings can be accessed but not reported during the attentional blink. *Nature* 383, 616–618.
- Marois, R., Chun, M.M., Gore, J.C., 2000. Neural correlates of the attentional blink. *Neuron* 28 (1), 299–308.
- Martens, S., Munneke, J., Smid, H., Johnson, A., 2006. Quick minds don't blink: electrophysiological correlates of individual differences in attentional selection. *J. Cogn. Neurosci.* 18 (9), 1423–1438.
- McLaughlin, E.N., Shore, D.I., Klein, R.M., 2001. The attentional blink is immune to masking-induced data limits. *Q. J. Exp. Psychol., A* 54 (1), 169–196.
- Nakayama, K., Mackeben, M., 1989. Sustained and transient components of focal visual attention. *Vision Res.* 29 (11), 1631–1647.
- Nieuwenhuis, S., Aston-Jones, G., Cohen, J.D., 2005a. Decision making, the P3, and the locus coeruleus–norepinephrine system. *Psychol. Bull.* 131 (4), 510–532.
- Nieuwenhuis, S., Gilzenrat, M.S., Holmes, B.D., Cohen, J.D., 2005b. The role of the locus coeruleus in mediating the attentional blink: a neurocomputational theory. *J. Exp. Psychol. Gen.* 134 (3), 291–307.
- Olivers, C.N., 2007. The time course of attention: it is better than we thought. *Curr. Dir. Psychol. Sci.* 16 (1), 11–15.
- Olivers, C.N., van der Stigchel, S., Hulleman, J., 2005. Spreading the sparing: against a limited-capacity account of the attentional blink. *Psychol. Res.* 8, 1–14.
- O'Reilly, R.C., Munakata, Y., 2000. *Computational Explorations in Cognitive Neuroscience: Understanding the Mind by Simulating the Brain*. MIT Press.
- Potter, M.C., 1993. Very short-term conceptual memory. *Mem. Cogn.* 21 (2), 156–161.
- Potter, M.C., Staub, A., O'Conner, D.H., 2002. The time course for competition for attention: attention is initially labile. *J. Exp. Psychol. Hum. Percept. Perform.* 28 (5), 1149–1162.
- Raymond, J.E., Shapiro, K.L., Arnell, K.M., 1992. Temporary suppression of visual processing in an RSVP task: an attentional blink? *J. Exp. Psychol. Hum. Percept. Perform.* 18 (3), 849–860.
- Rolke, B., Heil, M., Streb, J., Hennighausen, E., 2001. Missed prime words within the attentional blink evoke an N400 semantic priming effect. *Psychophysiology* 38 (2), 165–174.
- Seiffert, A.E., Di Lollo, V., 1997. Low-level masking in the attentional blink. *J. Exp. Psychol. Hum. Percept. Perform.* 23 (4), 1061–1073.
- Sessa, P., Luria, R., Verleger, R., Dell'Acqua, R., 2006. P3 latency shifts in the attentional blink: further evidence for second target processing postponement. *Brain Res.* 1137, 131–139.
- Shapiro, K.L., Caldwell, J., Sorensen, R.E., 1997a. Personal names and the attentional blink: a visual “cocktail party” effect. *J. Exp. Psychol. Hum. Percept. Perform.* 23 (2), 504–514.
- Shapiro, K.L., Driver, J., Ward, R., Sorensen, R.E., 1997b. Priming from the attentional blink: a failure to extract visual tokens but not visual types. *Psychol. Sci.* 8 (2), 95–102.
- Shea-Brown, E., Gilzenrat, M.S., Cohen, J.D., submitted for publication. Optimization of Decision Making in Multilayer Networks: the Role of Locus Coeruleus.
- Sperling, G., 1960. The information available in brief visual presentations. *Psychol. Monogr. Gen. Appl.* 74, 1–29.
- Usher, M., Cohen, J.D., Servan-Schreiber, D., Rajkowski, J., Aston-Jones, G., 1999. The role of locus coeruleus in the regulation of cognitive performance. *Science* 283 (5401), 549–554.
- Visser, T.A., Bischof, W.F., Di Lollo, V., 1999. Attentional switching spatial and non spatial domains: evidence from the attentional blink. *Psychol. Bull.* 125 (4), 458–469.
- Vogel, E.K., Luck, S.J., 2002. Delayed working memory consolidation during the attentional blink. *Psychon. Bull. Rev.* 9 (4), 739–743.
- Vogel, E.K., Luck, S.J., Shapiro, K.L., 1998. Electrophysiological evidence for a postperceptual locus of suppression during the attentional blink. *J. Exp. Psychol. Hum. Percept. Perform.* 24 (6), 1656–1674.
- Wyble, B., Bowman, H., 2005. Computational and experimental evaluation of the attentional blink: testing the Simultaneous

Type Serial Token model. In: Bara, B.G., Barsalou, L.W., Bucciarelli, M. (Eds.), CogSci 2005, XXVII Annual Conference of the Cognitive Science Society. Cognitive Science Society through Lawrence Erlbaum, pp. 2371–2376.

Wyble, B., Bowman, H., Nieuwenstein, M., submitted for

publication-a. A Computational Account of the Attentional Blink: Sparing, Blinking and Blanking.

Wyble, B., Bowman, H., Potter, M.C., submitted for publication-b. Categorically Defined Targets Trigger Spatiotemporal Attention.