The delayed consolidation hypothesis of all-or-none conscious perception during the attentional blink, applying the ST^2 framework

H. Bowman[1], P. Craston[1], S. Chennu[1], B. Wyble[2]

[1] Centre for Cognitive Neuroscience and Cognitive Systems, University of Kent, UK

[2] Department of Brain and Cognitive Sciences, MIT, USA

Abstract

There is evidence that conscious perception during the attentional blink is all-or-none, which contrasts with the finding of a continuum of perceptual strength in masking experiments. This seems to reveal the underlying representation of strength that can be found in the brain in the absence of attentional enhancement. We provide electrophysiological support for this all-or-none theory. Then, using principles from the ST^2 model of temporal attention and working memory, we show how this all-or-none pattern can be explained by the delayed target consolidation that is known to arise during the attentional blink.

Keywords: Attentional Blink, P300, ST^2

Introduction

Intuitively, it would seem that conscious perception is all-or-none. For example, bistable visual stimuli, e.g. the Necker cube or binocular rivalry, exhibit a dichotomy in perception: one or other stable percept is perceived in its entirety, without access to intermediate images. In other words, we either see *all* of a particular stable percept or *none* of it. Consistent with this perspective, (Sergent & Dehaene, 2004; Sergent, Baillet & Dehaene, 2005) have argued that conscious perception during the attentional blink (AB) is all-or-none; that is, targets presented during the attentional blink deficit are either completely perceived or completely missed, as revealed by subject visibility ratings. However, masking experiments in which strength is systematically varied suggest a continuous (although probably non-linear) gradation of visibility (Sergent & Dehaene, 2004, Del Cul, Baillet & Dehaene, 2007). This paper seeks to explain this seeming inconsistency in experimental findings, by arguing that all-or-none conscious perception during the attentional blink (AB) arises because of the delayed consolidation that is known to occur during the AB (Sessa et al., 2006; Vogel & Luck, 2002; Bowman et al., 2008). We do this by presenting a neural model of all-or-none consolidation during the AB, which reuses principles from the simultaneous type, serial token (ST^2) model of temporal attention (Bowman & Wyble, 2007; Wyble et al., 2009).

Why, though, should we be interested in the attentional blink? The AB is the deficit in reporting a second target (T2) if it is presented between 100 and 500 ms after a reported first target (T1), when a rapid serial visual presentation (RSVP) format is employed. The first reason is that the AB reflects a late stage deficit; that is, T2s that are missed during the AB are processed for visual and semantic features (Vogel et al., 1998). Thus, the AB seems to isolate the extra processing step that is required to encode an otherwise preconscious representation into working memory, i.e. to consciously perceive it. A second reason for being interested in the AB, is that it enables bottom-up strength and attention to be independently manipulated. A first requirement for an item to be consciously perceived is sufficient bottom-up strength. Under normal viewing conditions, most stimuli are strong enough to be consciously perceived. However, if stimulus representations are fleeting, as in RSVP, they will sometimes be too weak to enter consciousness. Hence, the weaker a neural representation of a stimulus, the less the likelihood it will enter awareness (Kanwisher, 2001).

However, this cannot be the sole requirement for conscious perception, as stimuli that are equally strong in perceptual terms, in some cases succeed but in others fail to enter consciousness (Luck et al., 1996, Rees et al., 2000). In particular, the neural representation of a stimulus also needs to be attended to before it can enter awareness. Koch & Tsuchiya (2007) have argued that attention and consciousness 'are distinct phenomena that need not occur together and can be manipulated using distinct paradigms.' To investigate the influence of target strength and attention on conscious perception, we thus require an experimental paradigm where the availability of attention and bottom-up strength can be manipulated independently (Kim & Blake, 2005); the AB is such a paradigm.

A key aspect of modern cognitive neuroscience is the identification of neurophysiological correlates of conscious perception. Within the context of the AB, the P3 component of the event-related potential (ERP) is viewed as such a correlate. A P3 component is only evoked by those targets in RSVP that can be correctly reported (e.g. Kranczioch et al., 2003). Target items that are missed do not evoke a P3. The P3 is thus generally seen as an EEG correlate of encoding items into working memory (Vogel et al., 1998) and, by the same logic, a number of studies have proposed that the P3 serves as an index of conscious perception (e.g. Sergent et al., 2005, Kranczioch et al., 2007). In line with these previous studies, we use the P3 component as a correlate of conscious perception for targets in RSVP.

We present EEG data investigating how bottom-up target strength and the availability of attention modulate conscious perception of targets in RSVP. We compare the EEG signatures (specifically, the P3 components) of targets presented outside and during the AB. We find that bottom up strength of targets affects the P3 for targets presented outside the AB, but does not influence the P3 evoked by

targets presented inside the AB. ST^2 (a prominent AB model) cannot account for these findings, thus we describe a new model that proposes two phases (not to be confused with the two-stage theory of Chun & Potter (1995)) of target perception in RSVP. Phase 1 is strongly sensitive to target strength, while Phase 2 is only weakly sensitive to such strength. We show how this *two-phase strength sensitivity theory* accounts for the experimental results presented and also all-or-none behavioural visibility ratings (Sergent & Dehaene, 2004) and P3s (Sergent et al., 2005) during the AB.

Experimental Findings

Experiment 1. We presented alphanumeric characters in black on a white background. Stimuli were in Arial font and had an average size of 2.1°x 3.4° visual angle. Participants viewed RSVP streams in which a single target was embedded into a continuous stream of distractors presented at fixation. The target for each trial was chosen at random from a list of 14 capital letters (B, C, D, E, F, G, J, K, L, P, R, T, U, V); distractors could be any digit except 1 or 0. The target item's position in the stream varied between 10 and 54. Items were presented at the unconventionally fast rate of approx. 20 items per second (item duration 47.1ms; no inter-stimulus interval) to ensure accuracy was not at ceiling in this relatively easy task. An RSVP stream consisted of 70 items (total stream length 3.3 seconds). The EEG sampling rate was 2000Hz (digitally reduced to 1000Hz at a later stage) and the data was digitally filtered at low-pass 85Hz and high-pass 0.5Hz during recording. 20 electrodes were placed at the following locations according to the 10/20 system: Fp1, Fp2, Fz, F3, F4, F7, F8, Cz, C3, C4, C7, C8, Pz, P3, P4, P7, P8, Oz, O1, O2, T7 and T8.

Experiment 2. Stimulus presentation was as per Experiment 1 except for a reduction in average stimulus size (1.03° x 0.69° visual angle) to ensure that the paradigm produced a reliable AB. RSVP streams were preceded by a fixation cross in the centre of the screen. After 400ms, the cross turned into an arrow indicating the side at which the two targets would be presented. After 200ms, two streams of digits were simultaneously presented at an equal distance of 2.6° visual angle to the left and right of fixation. The RSVP stream consisted of 35 items presented for 105.9ms each with no inter-stimulus interval. In a trial, T1 and T2 were selected from a list of 18 possible targets (A, B, C, D, E, F, G, H, J, K, L, N, P, R, T, U, V, Y); distractors could be any digit except 1 or 0. T1 appeared between position 5 and 17; T2 followed T1 at position 1 (no intervening distractors - lag 1), position 3 (2 intervening distractors - lag 3) or position 8 (7 intervening distractors - lag 8). The two-stream paradigm is irrelevant to the presented results, but was employed for a separate study.

Before the experiment started, participants were told to keep their eyes fixated on the centre of the screen, as trials with eye movements would be identified in the EOG and excluded from the analysis. Participants were told to direct their covert attention towards the indicated stream and search for the two target letters. EEG parameters were as per experiment 1.

Empirical Principles. We did not externally vary stimulus strength (for instance by manipulating contrast). However, Craston et al. (2009) suggest a relationship between the identity of a target letter and behavioural accuracy, as well as P3 size. The intrinsic stimulus characteristics (i.e. the shape of a particular letter) thus allow us to classify target letters as belonging either to an easy or a hard target set, which in turn provides us with an *indirect* measure of target strength. Target letters were classified based upon the data published in Bowman & Wyble (2007). Specifically, we ranked letters according to their average accuracy when used as targets in Bowman & Wyble (2007). Then we performed a median split, classifying the top half of letters as easy and the bottom half as hard.

The analyses presented are based on behavioural and EEG data from experiments 1 and 2. The data for the *single target* (target outside the AB) condition is taken from Experiment 1, whereas the *T2 following T1 at lag 3* (target inside the AB) is from Experiment 2. The continuous data is segmented by extracting a time window of -200ms to 1000ms for the single target condition, and a time window of -500ms to 1000ms for the *T2 following T1 at lag 3* condition. The single target data is baselined to the -200ms to 0ms window preceding target presentation and the data for *T2 following T1 at lag 3* is baselined to the -500ms to -300ms window with respect to target presentation (or the -200ms to 0ms period before the onset of the T1). The number of trials left in each condition after artifact rejection are specified in figure 1 (which differs from behavioural accuracy since it is calculated independent of artifact rejection). Full details of the experimental methods used can be found in Craston (2009).

Behavioural results. The prerequisite for the EEG analysis is a significant difference in behavioural accuracy between easy and hard letters for both targets presented outside and inside the AB. When analysing the behavioural data from Experiment 1, target letters T, K, U, V, L, D and G are categorised as easy, whereas E, C, B, P, F, J and R belong to the hard category. For Experiment 2, target letters T, K, U, V, L, D, G, N and H are categorised as easy, whereas E, C, B, P, F, J, R, Y and A belong to the hard category. The difference in the number of letters per category is because Experiment 2 contained four additional target letters. However, to re-emphasise, both subdivisions are inherited from Bowman & Wyble (2007) and are thus a priori.

If we apply this analysis to the *T2 following T1 at lag 3* condition, the accuracy scores for targets belonging to the easy and hard categories are 66% (SEM 4) and 46% (SEM 5), respectively. The difference is highly significant: $F(1,17) = 59.4$, $MSE < 0.01$, $p < 0.001$. The *T2 following T1 at lag 3* condition from Experiment 2 is thus used to investigate EEG processing inside the AB. In order to investigate the processing of targets outside the AB, we can employ either the *T1 with T2 presented at lag 8* condition from Experiment 2 or the single target in RSVP from Experiment 1. The easy/hard analysis for the *T1 with T2 presented at lag 8* condition from Experiment 2, however, does not meet our requirements, since the difference between easy and hard targets is only marginally significant; easy 87% (SEM 2) vs. hard 82% (SEM 3), $F(1,17) = 4.4$, $MSE < 0.01$, $p = 0.051$. This is likely to be due to ceiling effects, as T1 lag 8 accuracy is relatively high for both easy and hard letters.

The *single target* condition overcomes this problem, as we increased presentation rate to 50ms per item in Experiment 1 for accuracy to be below ceiling. Single targets show a highly significant effect of target difficulty, mean accuracy is 82% (SEM 4) for easy and 62% (SEM 4) for hard letters: $F(1,19) = 94.1$, $MSE < 0.01$, $p < 0.001$. Hence, we employ the single target condition from Experiment 1 to investigate EEG processing outside the AB.

EEG results. In order to perform a statistical analysis of our results, we extract the mean (unsubtracted) P3 size per subject for each accuracy-target difficulty combination for both targets outside and inside the AB.

Targets outside the AB. Figure 1.A shows that mean P3 size is influenced by our indirect measure of target strength

for correctly reported targets presented outside the AB. Targets in the Easy-Correct condition (8.9mV, SEM 0.9) have a significantly larger P3 than targets in the Hard-Correct condition (6.7mV, SEM 1.0): $F(1,19) = 28.2$, $MSE = 1.5$, $p < 0.001$. In addition, there is a significant difference in P3 size between the Easy-Incorrect and Hard Correct conditions: $F(1,19) = 7.2$, $MSE = 25.9$, $p = 0.016$. However, for the incorrectly reported targets, strength has no effect and the difference in P3 sizes between the Easy-Incorrect and the Hard-Incorrect conditions is not significant: $F(1,19) = 0.1$, $MSE = 49.8$, $p = 0.779$.

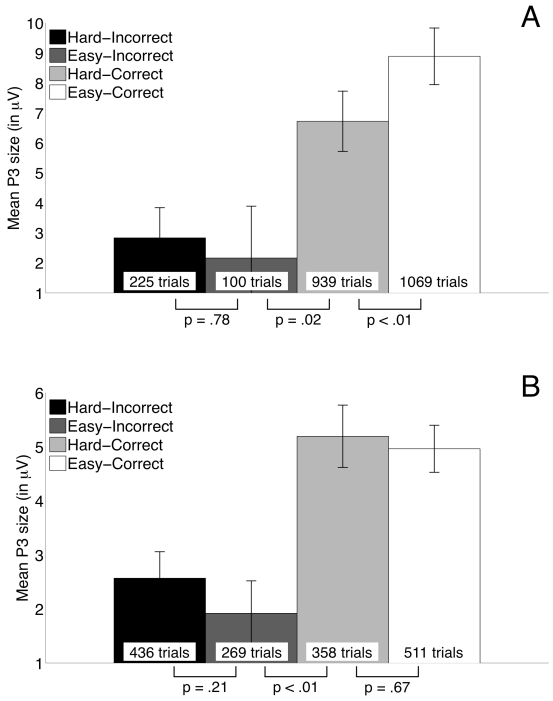

Figure 1: Bar charts displaying the mean P3 size (300-600ms with respect to target onset) for each accuracy-target difficulty combination. Error bars depict the standard error of the mean. Panel A: targets outside the AB. Panel B: targets inside the AB.

Targets inside the AB. Individual pairwise comparisons show that target difficulty does not have a significant effect on P3 size for targets inside the AB (see Figure 1.B). Both for correctly reported targets (Easy-Correct 5.0mV (SEM 0.4) vs. Hard-Correct 5.2mV (SEM 0.6); $F(1,17) = 0.2$, $MSE = 2.5$, $p = 0.664$) and incorrectly reported targets (Easy-Incorrect 1.9mV (SEM 0.6) vs. Hard-Incorrect 2.6mV (SEM 0.5); $F(1,17) = 1.7$, $MSE = 2.2$, $p = 0.209$) the difference between P3 sizes is not significant. The difference between the Easy-Incorrect and Hard-Correct conditions, however, is significant ($F(1,17) = 26.4$, $MSE = 3.7$, $p < 0.001$), which suggests that the high p-values in the previous analyses are not due to a lack of statistical power. Instead, this indicates that there is indeed no difference in P3 size between easy and hard targets both when they are correctly and incorrectly reported.

Discussion. In terms of behavioural accuracy, the strength of a target letter has an influence on target report both if the target is presented in isolation (i.e. outside the AB) and if it is presented during the AB. In Craston et al. (2009), we showed that target difficulty affects the size of the P3 for individually presented targets (i.e. outside the AB) that are correctly reported. We proposed that an easy target letter has more bottom-up strength than a hard letter and this increases the size of the P3 evoked by easy targets. In line with those findings, the results presented here illustrate how, for targets outside the AB, 'easy-hardness' affects the P3 if the target is correctly reported.

We performed the same analysis for targets presented during the AB, where we know that intrinsic stimulus characteristics (i.e. whether the target is easy or hard) affect behavioural report. However, 'easy-hardness' does not influence the P3 for targets during the AB. Figure 1.B illustrates that there is no significant difference in P3 size for easy and hard targets, both if the target is correctly or incorrectly reported. It seems that the P3 is influenced by different factors depending on whether a target is presented outside or inside the AB, with the later showing an all (easy or hard correct) or none (easy or hard incorrect) pattern (note: the fact that P3s for incorrect trials are not zero in Figure 1.B is probably because the T1-P3 bleeds into the early latencies of the T2-P3 and it is this overspill that is being measured).

In addition, the ST^2 model (Bowman & Wyble, 2007) (and recent revisions of the model, i.e. eST^2 (Wyble et al., 2009)), cannot explain this all-or-none pattern. This is because, whether inside or outside the AB, bottom-up strength modulates both behavioural performance and P3 size. Specifically, a stronger target, whether inside or outside the AB, always has a greater chance of being encoded into WM and has a larger activation trace.

Neural Modelling

We propose a new neural model, which, while still remaining faithful to central principles of the ST^2, can explain the all-or-none pattern during the AB. This model is built upon the following ST^2 principles.

Simultaneous type representation. ST^2 assumes a first stage of processing in which types are extracted. Types reflect all featural properties of items, e.g. in the context of the experiments considered here, the letter identity of the target and constituent visual features of that letter. Furthermore, multiple types can be active simultaneously during stage 1. In ST^2, the final layer of stage 1 (which is type-rich) is the task filtered layer (TFL). A task demand system ensures that only targets can be active at this layer.

Token-based working memory. It is assumed that the output layer of stage 1, which enables simultaneous activation of complete (task relevant) type representations, feeds into a second stage at which tokens are sequentially bound to active types. Tokens record instance specific/episodic information, such as, when an item occurred relative to other items. In ST^2, associating a token with a type *is* the process by which items are encoded into working memory (WM).

Transient attentional enhancement. Detection of a target initiates a transient attentional enhancer called the blaster. This mechanism transiently (for around 150 ms) amplifies all representations in later layers of stage 1.

Furthermore, this enhancement is a major contributor to the 'virtual' P3 generated by ST^2, the substantial part of which is observed at late layers of stage 1. In addition, such transient attentional enhancement amplifies representations sufficiently to initiate binding of active types to the active token, i.e. WM encoding.

Delayed consolidation. There is considerable evidence that the P3 component is delayed during the AB (Sessa et al., 2006; Vogel & Luck, 2002; Bowman et al., 2008). This effect is obtained in ST^2 since blaster firing is delayed during the AB (encoding of the first target into WM suppresses the blaster). Consequently, the model's analogue of the P3 and WM encoding/consolidation is also delayed.

The two-phase strength sensitivity theory

The ST^2 model cannot explain our experimental results. Consequently, we propose a modified theory: the *two-phase strength sensitivity theory*, which is based on theoretical concepts from ST^2 with some new principles, and uses activation equations described in Wyble et al. (2009).

We do not describe early stages of visual processing, but focus on later stages. Specifically, the task filtered layer (TFL) from ST^2 is expanded into two layers, see figure 2. These are the localized TFL (lTFL) and the globalized TFL (gTFL). The first of these is assumed to be locally prescribed in the ventral visual stream, and not to contribute to the P3 component, while the latter is assumed to reflect a more global activation (perhaps, akin to the global brain scale state in Dehaene et al. (2003)), and is assumed to be the major contributor to the P3 component.

It is important to note that, in ST^2, bottom up strength is determined by input strength. This input strength ranges from what we will call strong to weak. The easy/hard categories are then sampled from this complete range, such that, the easy/hard sub ranges overlap, with the former having a higher mean than the latter. We will talk in terms of the easy (respectively hard) range of values, which themselves contain a profile of weak and strong values.

lTFL. Focusing first on the lTFL, figure 3 depicts the activation traces for varying target strengths according to the two-phase theory. Figure 3 only shows traces that yield a correct target report at the gTFL. Indeed, the average of all correct traces is shown for each condition. Thus, the traces for targets outside the AB are lower in amplitude, since they generate an earlier blasting and consequently higher performance. This then ensures that a broader distribution of traces contribute to the average, thereby pulling it down. The key elements of lTFL target activation are as follows.

Phase 1 - strongly strength sensitive. Phase 1 is strongly sensitive to target strength, and targets with different strength values have different activation profiles.

Phase 2 - weakly strength sensitive. Phase 2 succeeds phase 1 and is only weakly sensitive to target strength, in the sense that all strong targets transition to the same activation level (i.e., the common attractor). Targets with enough activation strength during phase 1 will have entered a common attractor state by the time that phase 2 starts. This common attractor is reached because sufficiently active lTFL units saturate, thereby losing the differentiation across the range of strong targets. Furthermore, it is only strong traces (which saturate and enter the common attractor) that eventually get reported as a target during the AB. This is because it is only these strong traces that generate sufficient gTFL activation to be reported.

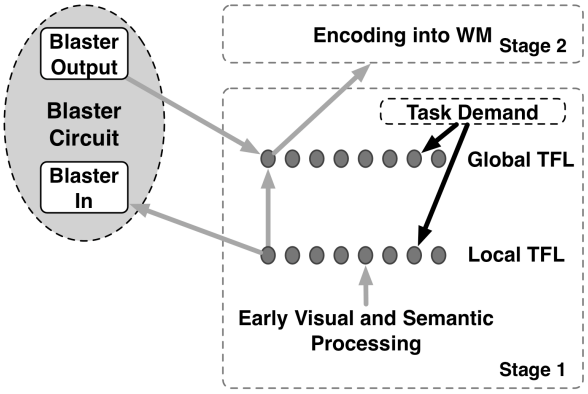

Figure 2: Revised ST^2 model implementing the two phase strength sensitivity hypothesis. The second stage, binding types to tokens using a binding pool, is not depicted, but is inherited unchanged from the basic ST^2 model (Bowman & Wyble, 2007). Links shown between individual pairs of units are replicated in one-to-one fashion across layers.

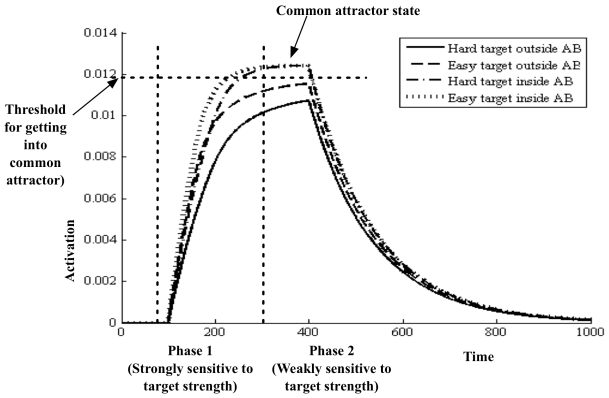

Figure 3: lTFL activation traces for seen targets with varying strengths. It is important to note that traces are plotted timelocked to the respective target onsets.

As initial support for this hypothesis, an ERP study by Del Cul et al. (2007) has indeed identified two phases of target processing, which have different sensitivities to target strength. Although not an RSVP study, Del Cul et al. (2007) manipulate target strength using masking. Importantly, they find an early phase that is highly sensitive to masking strength and a later phase, which - although not as weakly sensitive as we are proposing - is certainly a lot less sensitive than phase 1. This is most evident in figure 8 of Del Cul et al. (2007), in particular, the panel depicting ERP activity localised to posterior ventral temporal sources.

gTFL. We now discuss key aspects of the global TFL.

Attentional enhancement. The blaster provides attentional enhancement to targets. It fires once an item has been identified as a target. The enhancement increases activation

levels at the gTFL (see figure 4), which in turn initiates tokenisation (i.e. consolidation into WM) and the P3.

Behavioural accuracy. Only targets that generate sufficient gTFL activation are deemed to be reported; i.e., the area under a gTFL trace needs to cross a threshold, at which point the gTFL trace is terminated, and the model registers a correct report. This can be seen in figure 4.

Targets outside the AB. The hypothesised gTFL activation traces for targets presented outside the AB are the early traces in figure 4. As a target outside the AB is presented in isolation, the blaster is available to enhance the target's representation as soon as it is detected by the system. Thus, blaster enhancement occurs during phase 1, which is strongly sensitive to target strength. As seen in figure 4, the blaster increases target activation yielding a trace profile determined by the target's initial strength value. Weak targets (at the bottom of the easy or hard ranges), however, fail to fire the blaster. Consequently, they do not generate activation at the gTFL.

Phase 1 is critical for determining whether a target can be reported. Specifically, only Phase 1 targets over a particular threshold generate a sufficiently large gTFL trace to be reported. Thus, more targets in the easy category generate large enough gTFL traces to be reported than in the hard category. This direct relationship between target strength and the likelihood of target identification accounts for the easy-hard effect in behavioural accuracy for targets outside the AB: Hard targets have an accuracy of 66.7%, whereas easy targets have an accuracy of 100%.

The P3 is hypothesised to reflect the activation level of a target at the gTFL. However, the two-phase theory proposes that gTFL activation levels depend upon lTFL strength. Consequently, as seen in figure 4, outside the AB, the size of the P3 is determined by target strength. This accounts for the easy-hard effect on P3 size in our EEG results.

Furthermore, the two-phase theory can account for the results in Del Cul et al. (2007), who find that P3 size increases monotonically with target strength, as measured by the SOA between the target and the following mask ($p < 0.001$). It can be assumed that the shorter the SOA between the mask and the target, the more strongly the target's representation is weakened through masking. Shorter SOAs can thus be associated with lower target strength and, vice versa, the longer the SOA between target and mask, the higher the target strength. This is further evidence for target strength affecting P3 size when targets are presented in isolation (which, if we extrapolate to our experiment, would correspond to targets presented outside the AB).

Targets inside the AB. In line with ST^2, the two-phase theory suggests that the blaster is suppressed while T1 is encoded into WM. During the AB, T2 is presented before T1's WM encoding has completed. Consequently, the onset of the blaster is delayed for a target presented inside the AB and, as illustrated in figure 4, does not occur until phase 2, which is only weakly sensitive to target strength.

Weak targets (in the easy or hard ranges) do not have enough strength at lTFL to fire the blaster once it becomes available during phase 2. Hence, weak targets presented during the AB show no activation at the gTFL. Strong targets (within either the easy or hard ranges) presented during the AB, on the other hand, have enough strength to overcome the threshold for entering the lTFL common attractor during phase 1. These targets are in equal common attractor states when they fire the blaster. Hence, all targets that do manage to fire the blaster during the AB generate the same gTFL trace; see figure 4. The two-phase theory thus proposes that for targets presented during the AB, activation levels and thus the P3 are all-or-none.

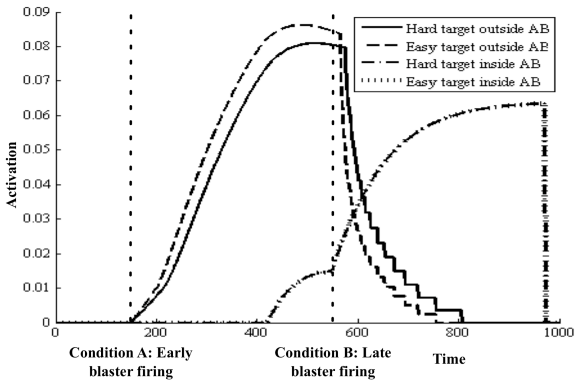

Figure 4: gTFL activation traces for seen targets. Condition A: Seen targets presented outside the AB. Condition B: Seen targets presented during the AB; Note that in Condition B, activation for the preceding target outside the AB is not shown, and that the traces for easy and hard targets overlap perfectly.

The probability of a target being reported is effectively determined in phase 1. For targets inside the AB, however, the blaster does not fire until after phase 1. Hence, behavioural accuracy is determined by a target's initial strength. Only strong targets overcome the threshold that allows entry into the common attractor. Thus, since strong targets are more frequent in the easy category, they are more likely to be reported than hard targets. Consequently, for targets inside the AB, there is an easy-hard effect on behavioural accuracy scores: Hard targets have an accuracy of 37%, whereas easy targets have an accuracy of 74%.

Conclusion

As discussed, there is behavioural evidence showing that participants' visibility is bimodal during the AB (Sergent & Dehaene, 2004). Observers were asked to report the extent to which the target had been perceived using a visibility scale ranging from 'Nothing' (0%) to 'Maximum visibility' (100%). For targets inside the AB, the majority of responses were concentrated around the minimum and maximum of the visibility scale. For targets presented outside the AB, however, the responses were gradually distributed with no clear threshold in visibility rankings. With respect to visibility, the P3 was also found to be distributed in an all-or-none fashion during the AB (Sergent et al., 2005). Trials with higher visibility scores showed a large P3, whereas trials with low visibility scores showed virtually no P3.

Like Sergent et al. (2005), we find that the P3 is unaffected by target difficulty (i.e. shows an all-or-none pattern) when targets are presented inside the AB. For targets outside the AB, we find that the P3 varies with target

difficulty, which is in agreement with the gradually distributed responses from Sergent et al. (2005). If we assume that visibility rating is governed by gTFL activation (which would be the natural interpretation), then the two-phase theory explains Sergent et al.'s findings.

Sergent et al. (2005) argue that conscious perception is all-or-none when attention is absent, as is the case during the AB. We argue though that it is not the absolute unavailability of attention that causes the all-or-none pattern. Rather, it is the delay of attentional enhancement that causes all-or-none during the AB (and consistent with this hypothesis, we know that T2 consolidation is indeed delayed during the AB (Vogel & Luck, 2002)). Specifically, the two-phase theory proposes that, due to the delayed firing of the blaster during the AB, enhancement of target activation does not have its effect until phase 2, which is only weakly sensitive to target strength. In phase 2, targets have either entered a common attractor at the lTFL, or their activation is too weak to get reported. Hence, activation traces show an all-or-none profile at the gTFL.

For targets outside the AB, however, the two-phase theory suggests that the blaster enhancement occurs during phase 1, which is sensitive to target strength. Hence, the strength of the percept varies and one would expect a graded continuum of conscious perception, as reported in Sergent et al. (2005). Further, we find a difference in behavioural accuracy scores between target letters belonging to easy and hard categories, both outside and inside the AB. This pattern of accuracy scores is obtained in the model since the likelihood that a target will be reported is determined by phase 1, which is strength sensitive.

It is important to consider the scope of the results presented here. In particular, the finding of an all-or-none pattern is specifically focused on the attentional blink window. Thus, we are not claiming to have provided evidence for all-or-none throughout conscious experience. Indeed, consistent with Sergent & Dehaene (2004), evidence for all-or-none was not found for isolated stimuli. Furthermore, Overgaard et al.'s (2006) critique of Sergent and Dehaene's work is focused on the broader all-or-none hypothesis and is thus largely avoided by the more restrictive claim considered here. In fact, our particular conjecture is that all-or-none patterns of conscious perception arise specifically when the delay between stimulus presentation and consolidation is long. Investigating all-or-none patterns in further experimental contexts where delayed consolidation may arise would be an important next research step.

References

Bowman, H., & Wyble, B. (2007). The Simultaneous Type, Serial Token Model of Temporal Attention and Working Memory. *Psychological Review*, 114 (1), 38–70.

Bowman, H., Wyble, B., Chennu, S., & Craston, P. (2008). A reciprocal relationship between bottom-up trace strength and the attentional blink bottleneck: Relating the LC-NE and ST^2 models. *Brain Research*, 1202, 25–42.

Craston, P. (2008). Applying Cognitive Electrophysiology To Neural Modelling Of The Attentional Blink. *PhD Thesis,* University of Kent.

Craston, P., Wyble, B., Chennu, S., & Bowman, H. (2009). The attentional blink reveals serial working memory encoding: Evidence from virtual & human event related potentials, *Journal of Cognitive Neuroscience*, 21 (3), 550–566

Dehaene, S., Sergent, C., & Changeux, J.-P. (2003). A Neuronal Network Model linking Subjective Reports and Objective Physiological Data during Conscious Perception. *Proc. Nat. Acad. Sci.*, 100 (14), 8520–8525.

Del Cul, A., Baillet, S., & Dehaene, S. (2007). Brain Dynamics Underlying the Nonlinear Threshold for Access to Consciousness. *PLoS Biology*, 5 (10), e260.

Kanwisher, N. (2001). Neural events and perceptual awareness. *Cognition*, 79 (1–2), 89–113.

Kim, C.-Y., & Blake, R. (2005). Psychophysical magic: rendering the visible 'invisible'. *Trends in Cognitive Sciences*, 9 (8), 381–388.

Koch, C., & Tsuchiya, N. (2007). Attention and consciousness: two distinct brain processes. *Trends in Cognitive Sciences*, 11 (1), 16–22.

Kranczioch, C., Debener, S., & Engel, A. (2003). Event-Related Potential Correlates of the Attentional Blink Phenomenon. *Cognitive Brain Research*, 17 (1), 177–187.

Kranczioch, C., Debener, S., Maye, A., & Engel, A. (2007). Temporal dynamics of access to consciousness in the attentional blink. *NeuroImage*, 37 (3), 947–955.

Luck, S., Vogel, E., & Shapiro, K. (1996). Word meanings can be accessed but not reported during the attentional blink. *Nature*, 383 (6601), 616–618.

Overgaard, M., Rote, J. Mouridsen, K. & Ramsøy, T.Z. (2006). Is conscious perception gradual or dichotomous? A comparison of report methodologies during a visual task. *Consciousness and Cognition*, 15, 700–708.

Rees, G., Wojciulik, E., Clarke, K., Husain, M., Frith, C., & Driver, J. (2000). Unconscious activation of visual cortex in the damaged right hemisphere of a parietal patient with extinction. *Brain*, 123 (8), 1624–1633.

Sergent, C., & Dehaene, S. (2004). Is Consciousness a Gradual Phenomenon? Evidence for an All-or-Nothing Bifurcation during the Attentional Blink. *Psychological Science*, 15 (11), 720–728.

Sergent, C., Baillet, S., & Dehaene, S. (2005). Timing of the brain events underlying access to consciousness during the attentional blink. *Nature Neuro.*, 10 (10), 1391–1400.

Sessa, P., Luria, R., Verleger, R. & Dell'Acqua (2006), R. P3 latency shifts in the attentional blink: Further evidence for second target processing postponement. *Brain Research*, 1137, 131–139

Vogel, E., & Luck, S. (2002). Delayed Working Memory Consolidation during the Attentional Blink. *Psychonomic Bulletin & Review*, 9 (4), 739–743.

Vogel, E., Luck, S., & Shapiro, K. (1998). Electrophysiological Evidence for a Post-perceptual Locus of Suppression During the Attentional Blink. *Journal of Experimental Psychology: Human Perception and Performance*, 24 (6), 1656–1674.

Wyble, B., Bowman, H., & Nieuwenstein, M. (2009). The attentional blink provides episodic distinctiveness: Sparing at a cost. *Journal of Experimental Psychology: Human Perception and Performance*, 35 (2), 324–337.