Neural dynamics of the attentional blink revealed by encoding orientation selectivity during rapid visual presentation

Matthew F. Tang[1,2], Lucy Ford[1], Ehsan Arabzadeh[2,3], James T. Enns[4,5], Troy A.W. Visser[5], & Jason B. Mattingley[1,2,6,7]

1. Queensland Brain Institute, The University of Queensland, Brisbane, QLD, Australia
2. Australian Research Council Centre of Excellence for Integrative Brain Function, Victoria, Australia
3. Eccles Institute of Neuroscience, John Curtin School of Medical Research, The Australian National University, Canberra, ACT, Australia
4. Department of Psychology, The University of British Columbia, Vancouver BC, Canada
5. School of Psychological Sciences, The University of Western Australia, Perth, WA, Australia
6. School of Psychology, The University of Queensland, Brisbane, QLD, Australia
7. Canadian Institute for Advanced Research (CIFAR)

Corresponding author

Matthew F. Tang

Queensland Brain Institute

The University of Queensland

St Lucia, QLD, Australia

Email: m.tang1@uq.edu.au

Abstract

The human brain is inherently limited in the information it can make consciously accessible. When people monitor a rapid stream of visual items for two targets, they can typically report the first, but not the second target, if these appear within 200-500 ms of each other, a phenomenon known as the attentional blink (AB). No extant theory has pinpointed the neural basis for the AB, partly because conventional neuroimaging approaches lack the temporal resolution to adequately characterise the neural activity elicited by each stimulus event in a rapid stream. Here we introduce a new approach that can identify the precise effect of the AB on behaviour and neural activity. Specifically, we employed multivariate pattern analysis to extract feature-selective information carried by randomly-oriented Gabors within a rapid serial stream. We show that orientation selectivity is enhanced for correctly reported targets, and suppressed when the same items are missed. By contrast, no such effects were apparent for irrelevant distractor items. Our findings point to a new theoretical account that involves both short-range and long-range temporal interactions between visual items competing for consciousness.

Introduction

Despite the remarkable processing capacity of the human brain, it is found wanting when undertaking multiple tasks concurrently, or when several goal-relevant stimuli must be dealt with in rapid succession. These limitations are particularly evident when individuals are required to execute responses to multiple stimuli under time pressure[1,2], or when they must select and report relevant target items that appear briefly and in rapid succession[3-5]. Elucidating the source of these limitations has been one of the most persistently difficult challenges in neuroscience and psychology. While the neural bases for these information processing limits are currently not fully understood, it is widely assumed that they are adaptive because they permit individuation of discrete objects and provide a mechanism by which selected sensory events can gain exclusive control over the motor systems responsible for goal-directed action.

Here we address a long-standing debate in the literature on the widely-studied 'attentional blink' (AB) bottleneck, where observers often fail to report the second of two target items (referred to as T2) when presented within 200-500 ms of the first target (T1) in a rapid stream of distractor stimuli[3-5]. While the AB clearly reflects a temporal limitation of visual attention, it has been difficult to characterise the neural basis of this bottleneck. Functional magnetic resonance imaging (fMRI) lacks the temporal resolution to accurately characterise neural activity associated with rapid serial visual presentation (RSVP) rates of 8-12 Hz, which are commonly used to elicit the AB[6,7]. Even electroencephalography (EEG), which has relatively good temporal resolution, results in smeared responses to items in an RSVP stream[8]. Furthermore, when data from fMRI or

63 EEG studies are analysed using conventional, mass-univariate approaches, they yield
64 only an overall neural response amplitude, but provide no information about how neural
65 responses represent featural information carried by single stimuli (e.g., their orientation).
66 Here we overcome these limitations by combining recently-developed
67 multivariate modelling techniques[9-16] with a novel RSVP task designed to determine the
68 neural basis for the temporal attentional bottleneck. Forward-encoding modelling allows
69 for the neural representation of stimulus-selective information to be quantified from
70 patterns of brain activity, using multivariate linear regression[9-16]. This approach allows
71 us to go beyond a simple description of the overall amount of neural activity associated
72 with target and distractor events, and instead to explicitly measure the neural
73 representation of visual featural information – in this case, orientation-selective
74 responses elicited by the grating stimuli – separately for each item within an entire
75 RSVP stream.
76 Using this approach, we set out to address two central theoretical questions.
77 First, does selection of a target from within an RSVP stream increase the gain or the
78 fidelity of its internal representation? Previous efforts to answer this question in the
79 domain of spatial attention have come from single cell recordings in non-human
80 primates[17,18], as well as whole-brain activity measured using fMRI[15,19] and EEG[9] in
81 humans. With few exceptions[20], these studies have found that spatial attention
82 increases the gain of feature-selective processing of attended stimuli. By contrast,
83 feature-based manipulations of attention, in which specific stimulus characteristics such
84 as colour or motion are cued for selective report, typically result in a sharpening of

neural selectivity[20-22]. To date, it remains unknown whether the limits of temporal attention in the AB are associated with changes in neural tuning associated with individual features of target and distractor items in the RSVP stream. Previous studies of activity in human primary visual cortex[6] and macaque lateral intraparietal area[23] have shown that neural responses to the second target are reduced overall on AB trials compared with non-AB trials, while subtraction-based EEG designs have shown that a late-stage component of the ERP (the P3) is reduced 200-400 ms after target presentation[8]. Critically, however, these measures fail to determine how the AB affects the manner in which the brain represents visual information, which could conceivably reflect a reduction in gain, an increase in tuning sharpness, or both.

A second unresolved theoretical question concerns the source of the temporal limitation. Existing theories have attributed the AB to the requirement to process the first target or to inadvertent distractor processing. In some theories, it is assumed that all items generate representations in early visual areas, but that the system inhibits items after T1 detection to avoid contamination by distractors[4,24-26]. On other accounts (so-called 'distractor-based' theories) the AB is assumed to reflect a cost associated with switching between target and distractor processing[27], or difficulty mapping these different stimulus types to memory tokens[28]. Finally, a third set of accounts argues that the representation of the second target becomes merged with either the first target or the distractors[29,30]. This class of theories is motivated by the finding that the perceived order of targets is often reversed (i.e., T2 is reported as appearing before T1).

Our novel RSVP task consisted of a stream of randomly oriented Gabors, with two higher-spatial frequency targets set amongst lower-spatial frequency distractors (Figure 1A and Movie 1). At the end of the RSVP stream, participants were asked to reproduce the orientations of the two targets (T1 and T2; Figure 2B). Critically, the orientation of each Gabor in the stream was uncorrelated with the orientation of all other items (Figure 1C), thus permitting the use of linear regression analyses to separately extract the influence of each item in the stream on neural activity measured by EEG, and on behavioural reports of the orientations of the two targets. These features allowed us to quantify for the first time the influence of both targets and distractors on participants' perceptual reports and on their associated neural representations.

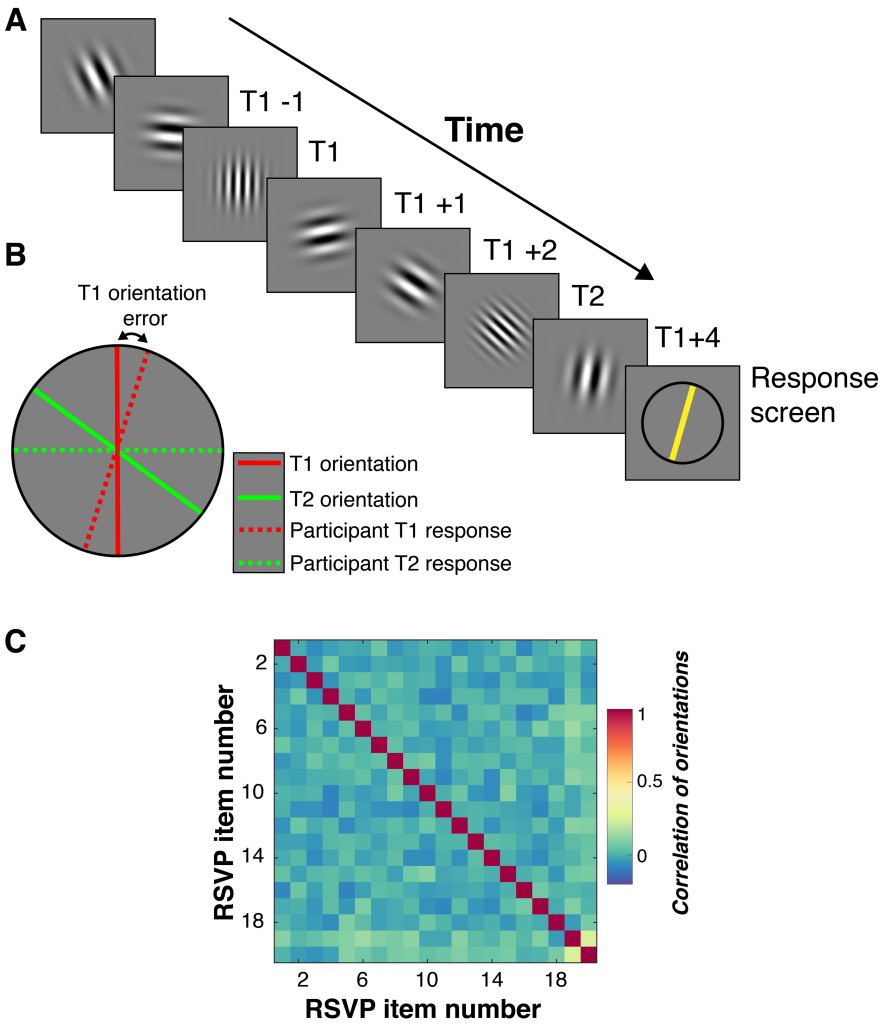

Figure 1. Rapid serial visual presentation (RSVP) task and illustration of correlations between Gabor orientations for measuring temporal visual selection. **A** Illustration of a typical trial in the RSVP task, which consisted of 20 sequentially presented Gabor patches at fixation. At the end of each RSVP stream, participants reproduced the orientation of both the first (T1) and second (T2) target stimuli (higher spatial frequency gratings). **B** Schematic of the feedback screen for responses. **C** Matrix of correlation values between orientations of the RSVP stimuli over trials in Experiment 1. As Gabor orientations were randomly drawn (without replacement) on each trial, across all trials the orientation of any given item in the stream was uncorrelated with the orientation of any other item. This permitted the use of regression-based approaches to isolate the behavioural and neural processing of individual items independently of surrounding items within the stream. The correlations were calculated for each participant and are displayed as averaged across participants.

Movie 1. Examples of two trials from the Gabor RSVP task.

To preview the results, an initial behavioural experiment replicated the key hallmarks of the AB effect. Most importantly, the orientation of the first target (T1) was reported with a relatively high degree of accuracy, whereas orientation judgements on the second target (T2) were degraded by guessing when T2 appeared 200-400 ms after T1. In a second experiment, we employed the same task while measuring EEG to examine neural activity associated with each item – both targets and distractors – within the RSVP stream. Forward encoding analyses of EEG activity showed that target items evoked greater orientation-selective information than distractor items when T2 was accurately reported (i.e., in non-AB trials), and that orientation information evoked by both targets was suppressed, relative to the distractors, when T2 was missed (i.e., in AB trials). Critical to our first question of whether focused attention influences the gain or fidelity of feature-specific stimulus representations, only the gain of the encoded EEG response was affected by T2 response accuracy. With respect to our second question — whether accuracy in registering the second target is linked to the processing of T1 or

148 to the intervening distractors — the evidence was in favour of T1-based theories of the
149 AB. We found no evidence to suggest that distractors influenced T2 processing. Finally,
150 we describe an entirely novel observation – one not predicted by any extant theory of
151 the AB – namely, a significant interaction between stimulus-specific feature
152 representations of T1 and T2, implying a previously unknown long-range temporal
153 integration of target representations within rapid sequential visual streams.

154 **Results**

155 **Experiment 1 - Behavioural measurement of target orientation perception for**
156 **stimuli in RSVP streams**

157 Twenty-two participants completed the RSVP task (Figure 1A). Each item within
158 a single RSVP trial was presented for 40 ms, with an 80 ms blank inter-stimulus interval
159 (ISI), thus yielding an 8.33 Hz presentation rate. The interval (or *lag*) between the first
160 (T1) and second (T2) targets was varied across trials to measure the temporal
161 dynamics of the AB. Participants reported the orientations of each of the two targets –
162 T1 and T2 – in order, at the end of the trial, by adjusting an on-screen cursor using a
163 mouse (Figure 1B). They were asked to determine the orientations as accurately as
164 possible, and were not given any time restriction. Visual feedback was provided at the
165 end of each trial.

166 ***Behavioural results replicate hallmarks of the attentional blink***

167 Participants' response errors (i.e., the difference between the presented and
168 reported orientation for each target) were centred around 0°, verifying that they were
169 able to perform the task as instructed. As shown in Figure 2A, accuracy was affected by

target position (T1 or T2) and Lag, a hallmark of the AB. Specifically, at Lag 1 accuracy for both T1 and T2 was degraded relative to accuracy at the other lags (2, 3, 5 & 7). Moreover, at Lags 2 and 3, T1 accuracy was high whereas T2 accuracy was relatively poor. Finally, at longer temporal separations (Lags 5 and 7) accuracy returned to a relatively high level for both target items.

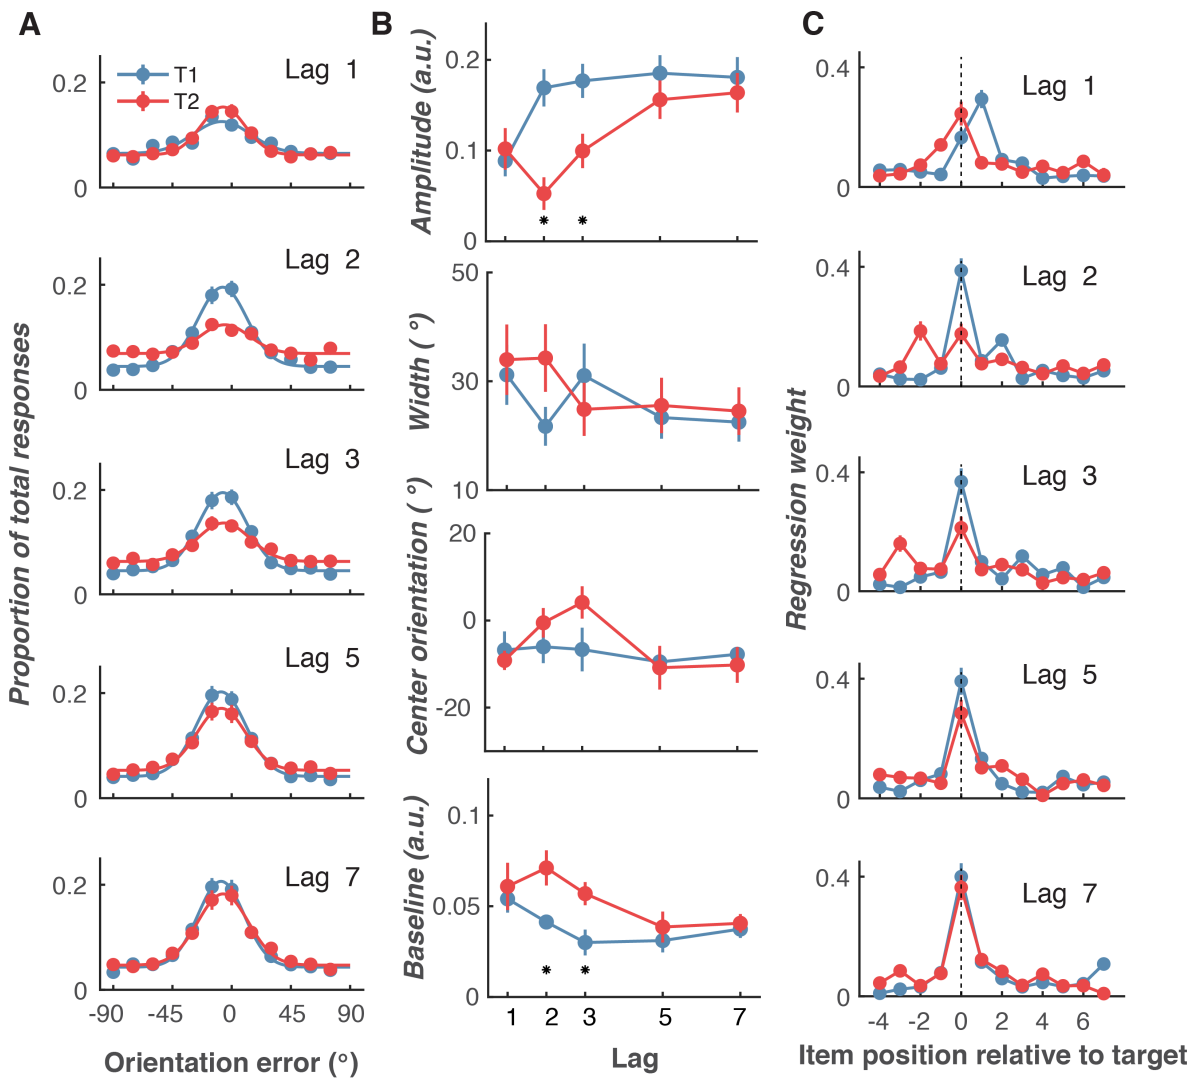

Figure 2. Behavioural results for RSVP task in Experiment 1. **A.** Aggregate response accuracy across participants (difference between presented and reported orientation) for the first (T1) and second (T2) target for each Lag condition. Line shows fitted Gaussian function. **B.** Quantified behavioural responses for the four parameters of the fitted Gaussian function. Amplitude gives the height, or gain, of the Gaussian, width shows the standard deviation, or fidelity, of the function. Centre orientation is the mean, which should be centred around 0 for unbiased estimates. Baseline is a constant parameter allowing for non-orientation selective guessing. Asterisks indicate $p < .05$. **C.** Regression results for influence of distractors and targets on participants' responses. Higher regression weights indicate that a given item was more influential for the reported orientation. Across all panels, error bars indicate ∓1 standard error of mean.

189 ***Gain or fidelity of behavioural response?***

190 We fitted Gaussian functions to each individual's data to characterise the nature
191 of the AB (Figure 2B; see Methods). The accuracy reduction for T2 at Lags 2 and 3 was
192 primarily linked to a reduction in amplitude/gain (all *ps*<0.0001). Critically for one of our
193 main questions, there was no effect of Lag on either the width (all *ps*>0.07) or precision
194 parameters (all *ps*>0.14). These results are consistent with a previous AB study using
195 similar analysis methods[31]. They also lend weight to the global workspace theory of
196 consciousness in the AB[32], which argues that participants either see the target and have
197 full awareness of it (allowing them to respond precisely), or they have no awareness
198 (and so simply guess randomly). By contrast, the results are inconsistent with the
199 opposing view that the AB involves a noisier (weaker fidelity) signal for the target that is
200 inaccurately reported[33].

201 ***Targets, not distractors, influence judgements of T1 and T2 orientation***

202 To evaluate the influence of distractors on participants' reports we aligned the
203 orientations of the items relative to target position within the RSVP stream (-4 to +7
204 items) and constructed a regression matrix to predict the behavioural response for each
205 target. If the orientation of an item is influential in determining what is reported, the
206 regression weight should be relatively high (Figure 2C). As expected, for all lags, each
207 target item was influenced principally by its own orientation. The one exception was the
208 stimulus at Lag 1, where perceptual reports of the orientation of T1 were also strongly
209 influenced by the orientation of T2. This observation is in line with numerous studies
210 which have suggested that temporal order information can be lost for consecutive

211 targets[29,30,34]. This phenomenon, also known as *Lag 1 switching,* where the perceived
212 order of the targets is reversed, explains why the accuracy of orientation judgements on
213 both T1 and T2 was reduced at Lag 1. By contrast, for items at Lags 2 and 3, orientation
214 judgements on T1 were only marginally influenced by the orientation of T2 (i.e., for
215 items at positions +2 and +3, respectively, in the RSVP stream). However, at these
216 same lags (where the AB is maximal) T2 reports were significantly influenced by T1
217 orientation (i.e., for items at positions -2 and -3, respectively). Importantly, there was no
218 reliable influence of distractors on reported target orientation at any lag, suggesting
219 distractors played little or no role in target orientation errors.

220 ***Long-range temporal integration of target orientations***

221 One account of the AB[29,30] has suggested that successive targets presented at
222 short lags are integrated into a single 'episodic trace', which accounts for Lag 1
223 switching, where the perceived order of the targets is reversed when they appear
224 sequentially. With the present task, we can directly quantify how targets are integrated
225 by looking for systematic biases in the reported orientation of a given target based on its
226 orientation difference with respect to the other target. Figure 3A shows orientation
227 judgement errors as a function of the difference between the two target orientations
228 presented on each trial. While the average orientation error is centred on 0°, confirming
229 that participants correctly reported the presented target, the perceived orientation of
230 either target (T1 or T2) was significantly biased toward the orientation of the other target
231 within the RSVP stream. Furthermore, these biases were orientation-tuned, such that
232 the largest bias occurred when targets differed by approximately 45°, somewhat

analogous to the recently-reported serial dependency effect[35-37]. The profile of biases suggests response integration, rather than replacement, as the latter would predict a large effect of T2 orientation on T1. On the contrary, but consistent with our linear regression analysis (see Figure 2C), there was only a relatively small influence of T2 orientation on participants' reports of T1. Instead, the bias reflected the *difference* between target orientations, which supports the idea that the critical features of the two targets are assimilated over time[29,30].

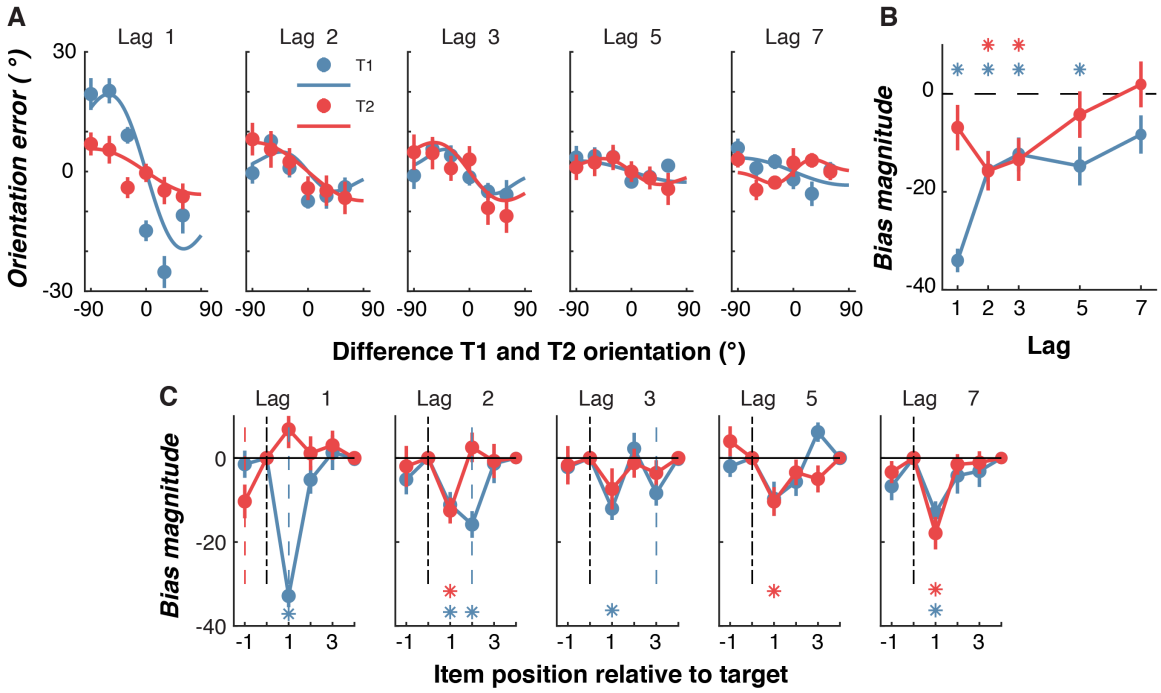

Figure 3. Influence of target and distractor stimuli on orientation judgements in Experiment 1. **A.** Orientation error (the difference between presented and reported orientation) plotted against the difference between presented T1 and T2 orientations (divided into 30° bins, for clarity of presentation). Negative values on the X-axis indicate that a given target was rotated clockwise relative to the other target. For instance, when examining T1, a negative value indicates that T2 was rotated clockwise relative to T1, whereas a positive value indicates that T2 was rotated anti-clockwise relative to T1. For T1, the plotted values reflect the calculation of T1 minus T2, and vice versa for the calculation of T2, to ensure values were equivalent for the comparison of interest. The same convention applies to orientation error, shown on the Y-axis. The fitted line is the first derivative of a Gaussian (D1) function showing the orientation-tuned amplitude and width of the response. **B.** Bias magnitude was quantified across participants by fitting the D1 function to each participant's (non-binned) data, with the amplitude showing bias magnitude. A negative amplitude on the Y-axis indicates that the perceived orientation was biased toward, rather than away from, the other target. **C.** Bias magnitude by difference with target and distractors. For both T1 and T2, the difference between the target and the item was found in the same manner as **A**. We fit the D1 function to find the magnitude of bias induced by each item. The dotted coloured lines indicate the temporal position of the other target (T1 = blue, T2 = red). For all panels, the asterisks indicate, for each target, at which lags the bias was significantly greater than zero (one-sample t-test $p<0.05$). Across all panels, error bars indicate ∓1 standard error of mean.

We fit first derivative of Gaussian (D1) functions[38-40] to quantify the amount of orientation-selective bias for both targets at each Lag for each participant. For T1 reporting, the bias was significantly greater than chance at all intervals except Lag 7, whereas for T2, there was a significant bias at Lags 2 and 3 only (Bonferroni corrected one-sample t-test, all *ps*<0.05). As might be expected[29,30], the 'attraction' bias in target reports was strongest when the two targets were presented with no intervening distractors between them (i.e., at Lag 1). An entirely novel finding, however, is that there was an equally strong attraction bias between targets presented at Lags 2 and 3 (see Figure 3B), even though participants were not explicitly aware of the orientation of T2 on AB trials.

Only the distractor immediately following each target biases reporting

Previous work suggests that distractor processing can significantly interfere with target processing[41-43], particularly for the immediate post-target item which can be integrated into the target representation[29,30,34]. To determine whether this was the case in our data, we repeated the previous analysis but used the difference between the target orientation and each of the items in the RSVP stream (Figure 3C). For most lags, the reported target orientation was significantly attracted toward the immediately following distractor, but was not reliably influenced by any other distractor. At Lag 1, there was no influence of distractors on reported orientations for either T1 or T2. This suggests that the representation of a given target stimulus is influenced by both the other target and by the post-target item. Taken together the results suggest that when the visual system detects a target, it automatically integrates features from the immediately subsequent stimulus.

Experiment 2 – Characterizing neural representations of target orientation information in RSVP streams

We next characterized the neural representation of individual items within each RSVP stream, and determined how these were affected by the AB. In Experiment 2, a group of 23 new participants undertook the RSVP task introduced in Experiment 1 while neural activity was concurrently measured using EEG. The paradigm was identical to the original in all respects, except that we now included targets at Lags 3 and 7 only (i.e., a single target inside and outside the AB, respectively) to increase within-subject power for the EEG analyses.

Behavioural results

The behavioural results were similar in all important respects to those found in Experiment 1. As shown in Figure 4A, participants performed well overall, and their orientation judgements for T1 and T2 were centred on the presented orientations. As in Experiment 1, T2 accuracy was significantly lower than T1 accuracy at Lag 3 ($t(22)=5.56$, $p<0.0001$), consistent with an AB effect, whereas there was no such difference at Lag 7 ($t(22)=0.90$, $p=0.38$; Figure 4B). Again, the behavioural effects at Lag 3 were due to an increase in the rate of guessing ($t(22)=2.07$, $p=0.05$) rather than a decrease in response precision ($t(22)=0.76$, $p=0.46$). Furthermore, as in Experiment 1, regression analysis showed that the reported orientation was driven mainly by the presented target (Figure 4C).

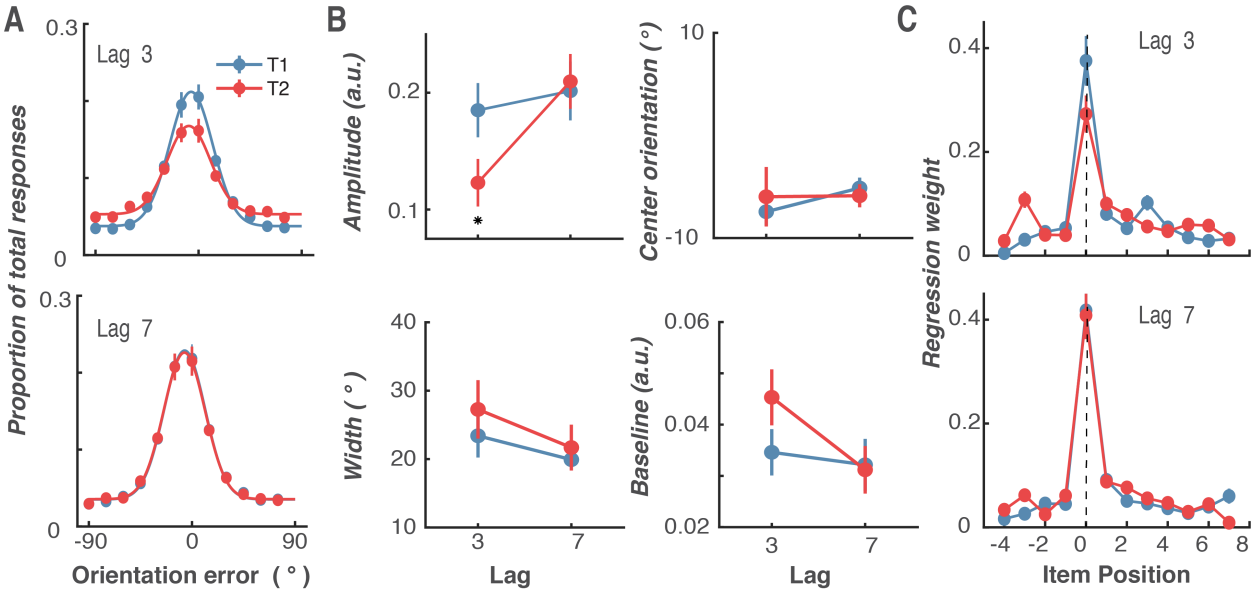

Figure 4. Behavioural results for the RSVP task in Experiment 2. **A.** Aggregate response accuracy across participants (difference between presented and reported orientations) for T1 and T2, shown separately for Lag 3 and Lag 7 trials. Lines are fitted Gaussian functions. **B.** Quantified behavioural responses for the four parameters of the fitted Gaussian functions (amplitude, width, centre orientation and baseline). Asterisks indicate $p<0.05$. **C.** Regression results for the influence of distractors and targets on the participants' responses. Error bars indicate ∓1 standard error of mean.

Orientation selectivity of RSVP items

We next applied forward modelling to the EEG data recorded during the RSVP task to quantify orientation information contained within multivariate patterns of neural activity. Because the orientations of successive stimuli within each RSVP stream were uncorrelated, we were able to quantify orientation selectivity for each grating without contamination from adjacent items within the stream. Forward encoding uses a linear regression-based approach to find multivariate patterns of EEG activity that are selective for features of interest – in this case orientation (see Methods for details). As no previous study has used forward encoding in conjunction with rapid visual presentations, we first verified that orientation selectivity for each of the 20 RSVP items

327 could be extracted separately using this approach. To do this, we constructed 20
328 encoding models, one for each of the items in the RSVP stream, based on the
329 orientations presented for that item across trials. Forward encoding was used to
330 determine, for each stimulus, whether there was an orientation-selective response
331 contained in the EEG activity, and at what time point any such response was evident.
332 As shown in Figure 5AB, the forward encoding analysis revealed robust and
333 reliable orientation selectivity contained within the pattern of EEG activity for each of the
334 grating stimuli presented during the RSVP. The presented orientation could be
335 successfully decoded over a time window that extended from 74 to 398 ms after the
336 item was presented. Examination of the neural responses to each of the 20 items within
337 the RSVP stream (Figure 5C) shows that orientation selectivity was evident as a series
338 of regularly spaced, short-lived impulse responses, each with a delay of around 50 ms
339 from grating onset and lasting approximately 300 ms. To quantify these observations,
340 we fit Gaussian distributions to the forward encoding results for each item separately for
341 each participant and at each time point. There was significant orientation selectivity
342 (given by the amplitude of the Gaussian) for each item immediately after it was
343 presented but not before (Figure 5D). These representations were temporally
344 overlapping, such that multiple orientation-selective responses (~3) were detectable at
345 the same time. Taken together, the forward encoding analyses verify that it is possible
346 to reliably recover the presented orientation of every RSVP item from the multivariate
347 pattern of neural activity recorded using EEG.

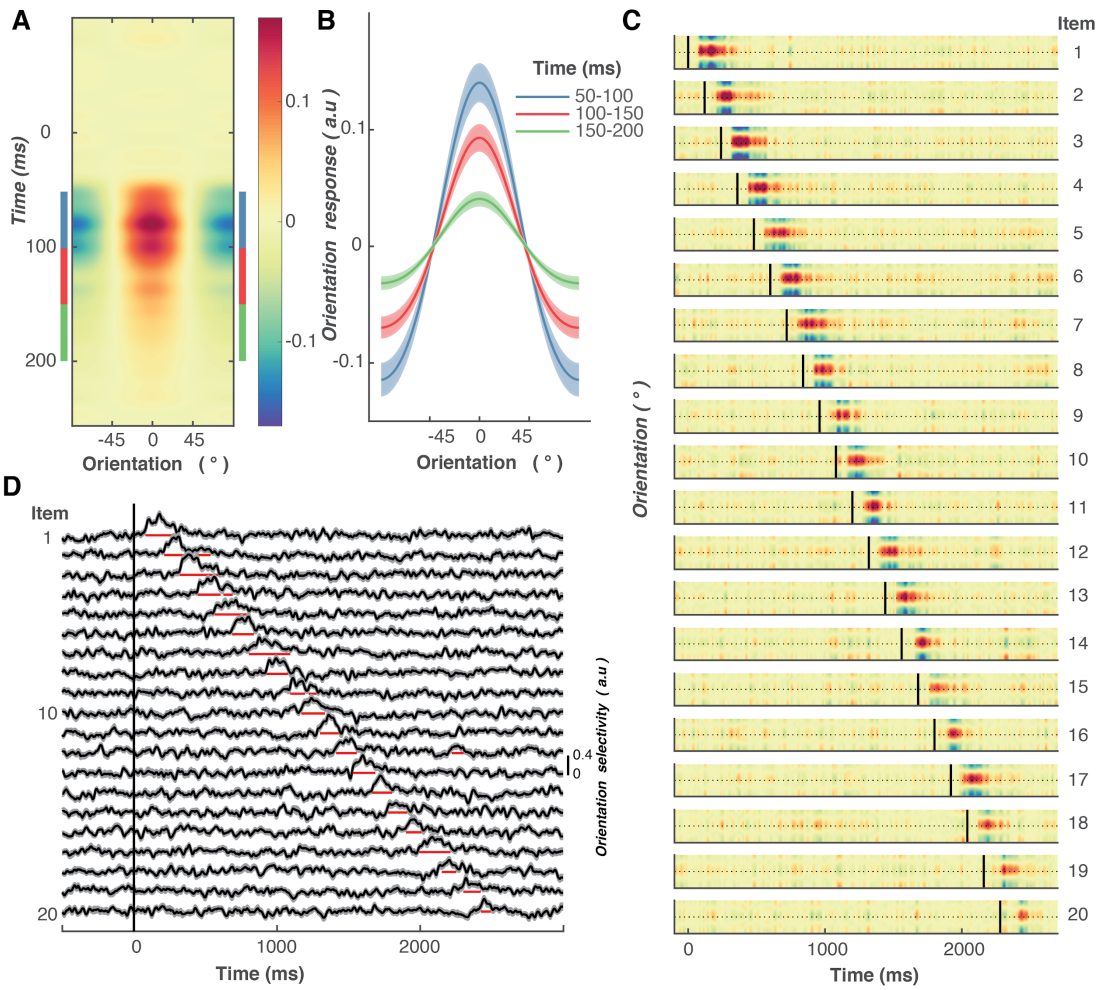

Figure 5. Reconstructing orientation selectivity for RSVP items using forward encoding modelling from EEG activity recorded in Experiment 2. **A**. Forward encoding results aligned at the time of stimulus onset and the presented orientation across all participants. All representations have been re-aligned so that the presented orientation is equivalent to 0°. **B.** Forward encoding results averaged over 50 ms bins (shown by corresponding colour in (A) following stimulus presentation. The orientation selectivity peaks around 50-120 ms after stimulus presentation. Error shading indicates ∓1 standard error of the mean across participants. **C.** Forward encoding results for each item in the RSVP stream. Vertical black lines indicate the presentation time of each of the 20 items within the RSVP stream. The colour scale is the same as in panel A. **D.** Gaussian distributions were fitted to each participant's data for each item in the stream, with the amplitude showing orientation selectivity. The red horizontal line segments underneath each trace indicate timepoints over which orientation selectivity was significantly different from zero (i.e., where orientation selectivity was greater than what would be expected by chance; two-tailed cluster-permutation, alpha $p<0.05$, cluster alpha $p<0.05$, N permutations = 20,000), which occurs immediately following item presentation. Grey shading in each trace indicates ∓1 standard error of the mean. a.u = arbitrary units.

The attentional blink is associated with reduced feature-specific information for the second target

We next examined how neural representations of the target items (T1 & T2) were influenced by the AB. To increase signal-to-noise for training the encoding model, we aligned the EEG data to the presentation time of each item in the RSVP stream and applied the same forward encoding procedure. This meant that the model was trained and tested across 12,000 presentations (600 trials by 20 RSVP items; see Figure 6). To determine the effect of the AB on orientation-selectivity, we separated the forward encoding results by target (T1, T2) and T2 accuracy (correct, incorrect). For the purposes of the analyses, trials were scored as correct if the reported orientation was within ±30 degrees of the presented orientation, a criterion which yielded roughly equal correct and incorrect trials at Lag 3. In line with the AB literature, we only included trials in which the participants correctly identified T1. Applying these criteria yielded a reliable AB effect, such that accuracy was significantly lower for T2 report versus T1 report at Lag 3 ($t(22)=3.79$, $p=0.001$), but not at Lag 7($t(22)=0.26$, $p=0.80$).

We again fitted Gaussians to each time point to quantify the amount of orientation-selective information evoked by the targets. For both T1 and T2, there was significant orientation-selective activity shortly after the relevant stimulus was presented (Figure 6). For Lags 3 and 7, there was no difference between correct and incorrect trials for the T1 representation. For T2, however, incorrect trials resulted in a significantly decreased orientation-selective response (cluster $p = .02$) than correct trials shortly after stimulus presentation (100 to 150 ms) at Lag 3. At Lag 7, there was no

significant difference in the orientation-selective response between correct and incorrect trials for T2. Performing the same analysis on the width parameter of the Gaussian showed no effect of the AB on the fidelity of orientation-selective information (Supplementary Figure 1).

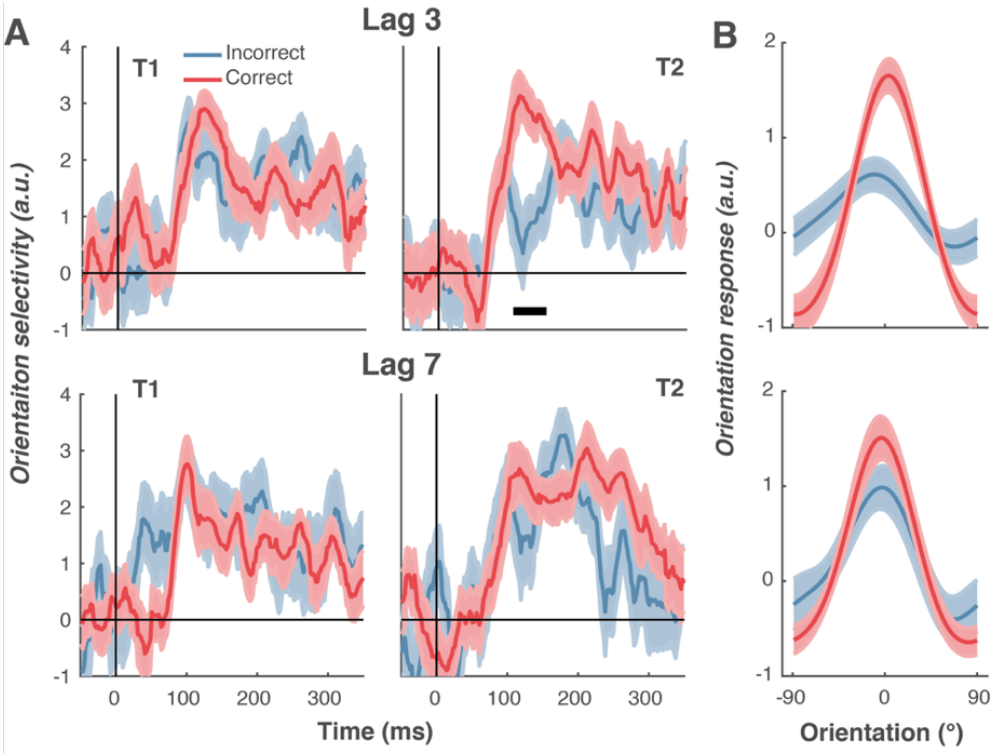

Figure 6. Neural representations of orientation-selective information during the attentional blink. **A.** Time course of measured orientation selectivity for T1 and T2, given by the amplitude of the fitted Gaussian parameter. Trials were scored as correct if the participant's response was within 30° of the presented orientation. Only trials in which participants responded accurately to T1 were included in the analysis. The thick black line in the upper right panel indicates a period of significant difference between Incorrect and Correct trials (two-tailed cluster-permutation, alpha $p<0.05$, cluster alpha $p<0.05$, N permutations = 20,000). **B.** Forward encoding results were averaged across the significant Lag 3 time points shown in A (upper right panel) to reconstruct a full representation of orientation information. Reliable orientation differences for T2 were present at Lag 3 (upper panel) but not at Lag 7 (lower panel). Shading indicates ∓1 standard error of the mean.

To determine how the AB affects the tuning of neural representations of target items, and to ensure we did not miss any small but consistent effects, we averaged the forward encoding results (Orientation × Time) over the earlier (100 to 150 ms) timepoints to increase signal-to noise-ratio, and recovered the orientation tuning curve (Figure 6B). Fitting Gaussians to these values confirmed that the AB was associated with a change in the gain of orientation selectivity for T2 at Lag 3, such that correct trials showed significantly greater amplitude in the tuning curve than incorrect trials ($t(22)=2.77$, $p=0.01$; Figure 6B upper panel). By contrast, the width of the representation was unaffected by the AB ($t(22)=0.90$, $p=0.38$) for the same item. There was no significant change in orientation selectivity for T2 at Lag 7 ($t(22)=0.89$, $p=0.38$; Figure 6B lower panel). The reduction in T2 selectivity for incorrect trials at Lag 3 was not driven by our arbitrary split of trials into correct and incorrect categories. To show this, we sorted the evoked T2 forward encoding results by the amount of orientation error. There was significantly greater orientation selectivity when the actual and reported orientations were similar, and this tuning effect decreased gradually with larger errors (Geissler-Greenhouse one-sample ANOVA, $F(1,22)=2.60$, $p=0.028$, Supplementary Figure 2). Taken together, these results indicate that the AB is associated with a reduction in gain, but not fidelity, of feature-selective information for the second target item (T2), and that this effect occurs soon after the target appears within the RSVP stream.

How do neural representations of targets and distractors contribute to the AB?

428 We next examined how targets and distractors are represented neurally in order
429 to directly test the different predictions made by T1- versus distractor-based accounts of
430 the AB. T1-based accounts[25,26] argue that the second target deficit is caused by
431 extended processing of the first target, whereas distractor-based accounts[27,28] instead
432 argue that deleterious processing of the distractors, mainly between T1 and T2, causes
433 the second target to be missed. The theories thus make distinct predictions about the
434 neural representation of target and distractor items. According to T1-based accounts,
435 target representations should be enhanced relative to those of distractors, and missed
436 T2 items on AB trials should be more poorly represented than correctly reported T2
437 items. By contrast, distractor-based accounts predict that neural representations of
438 distractor items should be stronger on AB trials than on non-AB trials.
439 As before, we averaged the forward encoding modelling results (Orientation ×
440 Time) across an early time point (100 to 150 ms), and fit Gaussians to each participant's
441 data to quantify orientation selectivity (Figure 7A). For correct trials (i.e., for trials in
442 which participants correctly reported the T2 orientation to within 30°), the two targets
443 resulted in significantly higher orientation selectivity (gain) than the immediately
444 adjacent distractors (-2,-1,+1 and +2 items) for both T1 and T2 representations (all
445 $ps<0.032$). On incorrect trials, orientation selectivity for T1 was not significantly greater
446 than selectivity for the surrounding distractors ($t(22)=0.25$, $p=0.80$), even though only
447 trials in which T1 was correctly responded to were included. Most interestingly, on
448 incorrect trials orientation tuning associated with T2 items was significantly lower than
449 orientation tuning for the immediately adjacent distractors ($t(22)=2.13$, $p=0.04$),

suggesting that the featural information carried by T2 stimuli was suppressed. Taken together, these results suggest that on trials in which participants accurately report target orientation the neural representations of target stimuli are boosted relative to those of the distractors. By contrast, missing the second target, as occurs during the AB, is associated with a significant *suppression* of featural information carried by the target.

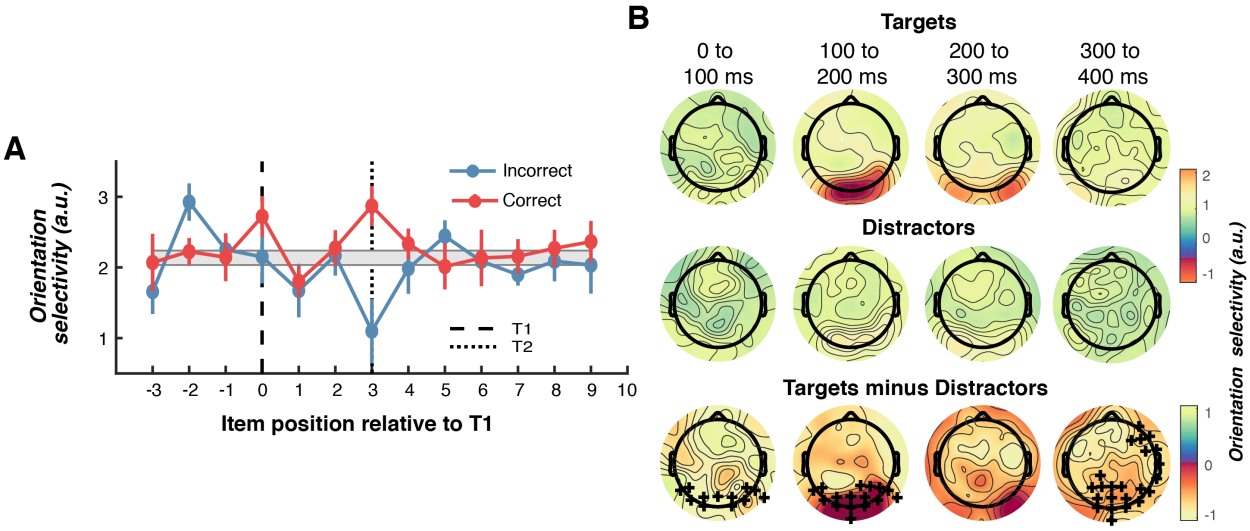

Figure 7. Orientation selectivity and scalp topographies for targets and distractors in Experiment 2. **A.** Neural orientation selectivity (amplitude of Gaussian) of target and distractor representations for Lag 3 trials. To aid comparison, the grey bar indicates the average distractor representation. Error bars indicate ∓1 standard error of mean. **B.** Headmaps showing univariate orientation selectivity over time, plotted separately for targets and distractors. Symbols indicate cluster-permuted significance comparing targets and distractors (two-tailed cluster-permutation, alpha $p<0.05$, cluster alpha $p<0.025$, N permutations = 1,500).

Localisation of orientation selectivity of targets and distractors

In a final step, to localise the reported neural effects, we performed a univariate sensor-level analysis for orientation selectivity[10] to find the topographies associated with target and distractor processing. To do this, we separately trained a simplified model of

orientation selectivity on each type of stimulus (targets and distractors) and separately for each EEG sensor. Orientation information for both targets and distractors was evident most strongly over occipital and parietal areas, and target items generated significantly greater selectivity over these areas than distractors (Figure 7B). These findings suggest that while target and distractor items are processed in overlapping brain regions, targets generate significantly greater orientation-selective information than distractors.

Discussion

We combined forward encoding modelling of EEG data, recorded as participants performed a novel AB task, in order to determine the neural and behavioural bases of the limits of temporal attention. The behavioural results replicated all the hallmarks of the AB. First, response accuracy was significantly reduced when T2 was presented within 200-400 ms of T1. Results from Experiment 2 revealed that reporting a target successfully depends on boosting its neural representation relative to other items in the RSVP stream. Second, missing the T2 item corresponded to a suppression of the second target representation relative to the distractors. These two findings positioned us to apply modelling on a trial-by-trial basis for both behavioural and EEG data. The results were decisive in showing that the gain, but not the fidelity, of a target's neural representation was influenced by the AB. We found that target representations clearly influenced one another, such that the reported orientation of one target was biased toward the orientation of the other. By contrast, there was no evidence for suppression of the neural representations of the distractors. This observation supports theories that

493 have attributed the second-target deficit to necessary first target processing[4,24], but is
494 inconsistent with theories that attribute the AB to inadvertent processing of distractor
495 items[27,28].

496 An important but unexpected result of Experiment 2 was that target reports were
497 influenced by one another despite being separated by several hundred milliseconds and
498 multiple distractor items. One influential theory assumes that the AB is caused by
499 temporal integration of the target stimulus with the immediate post-target distractor[29,30].
500 Our RSVP task revealed an unexpected but novel finding: target representations appear
501 to be influenced by each other, even when they are separated by multiple distractor
502 items. This finding is not explicitly predicted by any existing account of the AB. The
503 largest bias was for Lag 1 trials, in which the two targets appear sequentially, a result
504 that is consistent with Lag 1 switching[29,30,34]. The orientation of the immediate post-
505 target distractor also significantly biased perceived target orientation, whereas the
506 distractors that appeared between the targets did not bias participants' perceptual
507 judgements. Taken together, our findings across two experiments suggest that in RSVP
508 tasks the appearance of a target stimulus starts a period of local integration which
509 involuntarily captures the next item, whether it is a target or a distractor. This is followed
510 by a more global integration of any targets that have been selected, possibly within
511 working memory.

512 Our use of forward encoding modelling of the EEG data add to previous
513 behavioural results[31] by demonstrating that the gain in neural representations of Lag 3
514 stimuli is significantly reduced in trials in which an AB occurs, compared with trials in

515 which the T2 item is reported correctly. Supporting the behavioural results, there was no
516 effect on the fidelity (tuning width) of EEG-derived orientation selectivity during the AB.
517 The neural results also go beyond the behavioural findings by showing that the gain of
518 Lag 3 items is only reduced on AB trials, but is unaffected on non-AB trials. Taken
519 together, these results suggest that temporal attention operates in a similar manner to
520 spatial[15,17-19] attention, but not to feature-based attention[20,21], as the former has been
521 found to affect the gain of neural responses whereas the latter tends to affect the
522 sharpness of neural tuning.

523 The second major aim of our study was to resolve the persistent debate between
524 T1- and distractor-based theories of the AB [4,24-28]. Behaviourally, we found scant
525 evidence that distractors influence target perception. Consistent with T1-based
526 accounts of the AB, there were robust neural representations of distractor stimuli, but
527 there was no evidence that distractor representations were boosted following initial
528 target detection, as would be predicted by distractor-based accounts. Furthermore, we
529 found no evidence that post-T1 distractors were suppressed, as would be predicted by
530 T1-based inhibition accounts of the AB[4,24]. Instead, in line with the T1-based bottleneck
531 account, the representations of *both* targets were boosted relative to those of the
532 distractors. If the second target was missed, however – as occurs during the AB – then
533 the representation of the second target was significantly suppressed relative to the
534 distractors. Taken together, these results suggest that when the first target is processed
535 rapidly, attention is efficiently redeployed to the second target, causing its
536 representation to be boosted. By contrast, if the second target appears while processing

537 of the first target is ongoing, the visual system actively suppresses the information to
538 avoid the targets interfering with each other. There was no evidence that distractors
539 were suppressed for target processing. Overall, these results are most consistent with
540 T1-based accounts of the AB.

541 In summary, the current work adds to our understanding of the neural and
542 behavioural basis of temporal attention. For the first time, we were able to recover a
543 neural signature for each and every item within the RSVP stream, something that has
544 not been possible with conventional analytic approaches to EEG and fMRI data . Our
545 novel RSVP task indicated that while there is co-modulation of featural information
546 carried by each of the target stimuli, there is no evidence for distractor suppression in
547 such RSVP tasks. We also document for the first time the existence of long-range
548 integration of targets separated by several hundred milliseconds. The current work was
549 not designed to pinpoint the neural locus of these long-range effects, but future studies
550 could combine our behavioural task and analytic approach with fMRI in an effort to
551 localise their cortical origins. Some researchers have claimed that feedback and
552 feedforward processes modulate different aspects of the AB[44]. Again, our study was
553 not designed to address this issue, but it could potentially be resolved if our approach
554 were combined with calcium imaging or multi-electrode recordings in animal models,
555 particularly if multiple areas were recorded from simultaneously[45].
556

Online methods

Participants

In Experiment 1, 22 participants (13 females, 9 males; median age 22 years; range 19-33 years) were recruited from a paid participant pool and reimbursed at AUD$20/hr. In Experiment 2, 23 participants (14 females, 9 males; median age 23 years; range 19-33 years old) were recruited from the same pool and reimbursed at the same rate. Each person provided written informed consent prior to participation and had normal or corrected-to-normal vision. The study was approved by The University of Queensland Human Research Ethics Committee and was in accordance with the Declaration of Helsinki.

Experimental setup

Both experiments were conducted inside dimly-illuminated rooms. The stimuli were displayed on a 22-inch LED monitor (resolution 1920 x 1080 pixels, refresh rate 100 Hz) using the PsychToolbox presentation software[46,47] for MATLAB. In Experiment 1, participants were seated at a distance of approximately 45 cm from the monitor. In Experiment 2, the same viewing distance was maintained using a chinrest to minimise head motion artefacts in the EEG. At a viewing distance of 45 cm, the monitor subtended 61.18º x 36.87º (one pixel = 2.4’ x 2.4’).

Task

A schematic of the task is shown in Figure 1. Movie 1 shows two example trials. Each trial began with a central fixation point and the RSVP stream commenced after 300 ms. The stream consisted of 20 Gabors (0.71° standard deviation, ~5° diameter,

579 100% contrast, centred at fixation) on a mid-grey background. On each trial, the
580 orientations of the twenty Gabors in the stream were drawn pseudo-randomly without
581 replacement from integer values ranging from 0-179°. Each item was presented for 40
582 ms and was separated from the next item by a blank inter-stimulus interval of 80 ms,
583 yielding an 8.33 Hz presentation rate. The participants' task was to reproduce the
584 orientations of the two high-spatial frequency Gabors (targets; 2 c/°) while ignoring the
585 items of a low-spatial frequency (distractors; 1 c/°). Between 4 and 8 distractors, varied
586 pseudo-randomly on each trial, were presented before the first target (T1) to minimise
587 the development of strong temporal expectations, which can reduce the AB[42,48]. The
588 number of distractor items between T1 and T2 defined the inter-target lag (1, 2, 3, 5, 7
589 in Experiment 1, and 3 and 7 in Experiment 2).

590 Participants were asked to monitor the central RSVP stream until the
591 presentation of the last Gabor, after which a response screen appeared (see Figure
592 1B). The response screen consisted of a centrally-presented black circle (10° diameter)
593 and a yellow line. Participants rotated the line using a computer mouse to match the
594 perceived orientation of the target stimuli, and clicked to indicate their desired response.
595 They were asked to reproduce the orientations of the two targets (T1, T2) in the order
596 they were presented, and to respond as accurately as possible, with no time limit. After
597 providing their responses, participants were shown a feedback screen which displayed
598 their orientation judgements for T1 and T2, and the actual orientations of both targets
599 (see Figure 1C). The feedback was displayed for 500 ms before the next trial began,
600 and participants were given a self-paced rest break every 40 trials. There were 600

trials in each of the two experiments, with an equal distribution of trials across the lag conditions. Each experiment took between 50 and 60 minutes to complete.

EEG acquisition and pre-processing

In Experiment 2, continuous EEG data were recorded using a BioSemi Active Two system (BioSemi, Amsterdam, Netherlands). The signal was digitised at 1024 Hz sampling rate with a 24-bit A/D conversion. The 64 active scalp Ag/AgCl electrodes were arranged according to the international standard 10–20 system for electrode placement[49] using a nylon head cap. As per BioSemi system design, the common mode sense and driven right leg electrodes served as the ground, and all scalp electrodes were referenced to the common mode sense during recording. Pairs of flat Ag-AgCl electro-oculographic electrodes were placed on the outside of both eyes, and above and below the left eye, to record horizontal and vertical eye movements, respectively.

Offline EEG pre-processing was performed using EEGLAB in accordance with best practice procedures[50,51]. The data were initially down sampled to 512 Hz and subjected to a 0.5 Hz high-pass filter to remove slow baseline drifts. Electrical line noise was removed using the *clean_line.m,* and *clean_rawdata.m* in EEGLAB[52] was used to remove bad channels (identified using Artifact Subspace Reconstruction), which were then interpolated from the neighbouring electrodes. Data were then re-referenced to the common average before being epoched into segments for each trial (-0.5 s to 3.0 s relative to the first Gabor in the RSVP). Systematic artefacts from eye blinks, movements and muscle activity were identified using semi-automated procedures in the SASICA toolbox[53] and regressed out of the signal. The data were then baseline

corrected to the mean average EEG activity from 500 to 0 ms before the first Gabor in the trial.

Behavioural analysis

To determine how the AB affected participants' perception of target stimuli, for each trial we found the difference between the actual target orientation and the reported orientation (i.e., the *orientation error*) for T1 and T2. This approach is analogous to one employed in previous work that examined whether the AB is associated with discrete or graded awareness of the T2 stimulus[31]. For each lag condition, we found the proportion of responses (in 15° bins) between -90° and +90° for the orientation errors (see Figure 2A and 3A) and fit Gaussian functions with a constant offset (Equation 1) using non-linear least square regression to quantify these results for each participant (Figure 2B and 3B):

$$G(x) = A\, exp(-\frac{(x-\varphi)^2}{2\sigma^2}) + C \quad (1)$$

Where *A* is the amplitude, reflecting the number of responses around the reported orientation, φ is the orientation on which the function is centred (in degrees), σ is the standard deviation (degrees), which provides an index of the precision of participants' responses, and *C* is a constant used to account for changes in the guessing rate.

We used a regression-based (Figure 2D and 4C) approach to determine how targets and distractors within each RSVP stream influenced behavioural responses[54]. To do this, we aligned orientations of both distractor and target Gabors from 4 items prior to the appearance of T1 through to 9 items after the appearance of T2 to construct a regression matrix. Standard linear regression was used to then determine how the

645 orientations of the items affected the reported orientation for both T1 and T2, with a
646 higher regression weight indicating a more influential item.

647 To determine whether the finding that the orientations of T1 and T2 influenced
648 the reported orientation was due to participants integrating the other target or the
649 surrounding distractors (Figure 3), we found the difference in orientation between the
650 target of interest and the other item (either target or distractor) and the orientation error
651 for each trial. This showed an orientation-tuned effect characteristic of integration. To
652 quantitatively determine the magnitude of this effect, we fit first-derivative Gaussian
653 functions (D1; Equation 2) to these responses[39,40].

654 $$\mathrm{D1}(x) = A \times \frac{1}{\sigma} \times x - \varphi \times exp(-\frac{x-\varphi^2}{2\sigma^2}) \quad (2)$$

655 Where *A* is the amplitude, φ is the orientation on which the function is centred (in
656 degrees), and σ is the standard deviation (degrees).

657 **Forward encoding modelling**

658 Forward encoding modelling was used to recover orientation-selective responses
659 from the pattern of EEG activity for both target and distractor items in the RSVP stream.
660 The forward encoding technique has been used previously to reconstruct colour[16],
661 spatial[15] and orientation[19] selectivity from timeseries data acquired through functional
662 magnetic resonance imaging (fMRI). More recently, the same approach has been used
663 to encode orientation[9-13] and spatial[14] information contained within MEG and EEG data,
664 which have better temporal resolution than fMRI.

665 We used the orientations of the epoched data segments to construct a
666 regression matrix with 9 regression coefficients, one for each of the orientations. This

667 regression matrix was convolved with a tuned set of nine basis functions (half cosine
668 functions raised to the eighth power) centred from 0° to 160° in 20° steps. This tuned
669 regression matrix was used to measure orientation information either across trials or in
670 epoched segments. This was done by solving the linear equation (3):

671 $$B_1 = WC_1 \quad (3)$$

672 Where B_1 (64 sensors × N training trials) is the electrode data for the training set, C_1 (9
673 channels × N training trials) is the tuned channel response across the training trials, and
674 W is the weight matrix for the sensors to be estimated (64 sensors x 9 channels).
675 Following methods recently introduced for M/EEG analysis, we separately estimated the
676 weights associated with each channel individually[55]. W was estimated using least
677 square regression to solve equation (4):

678 $$W = (C_1\, C_1^T)^{-1}\, C_1^T\, B_1 \quad (4)$$

679 Following from previous work[11,55], we sought to remove any correlations between
680 sensors, as these prevent finding the true solution to the linear equation. To do this, we
681 first estimated the noise correlation between electrodes and removed this component
682 through regularisation[56]. The channel response in the test set C_2 (9 channels × N test
683 trials) was estimated using the weights in (3) and applied to activity in B_2 (64 sensors ×
684 N test trials), as per Equation 5:

685 $$C^2 = (W\, W^T)\, W^T\, B^2 \quad (5)$$

686 To avoid overfitting, we used cross validation (10-fold in the initial whole-trial analysis,
687 and 30-fold when the stimulus presentations were stacked), where X-1 of epochs were
688 used to train the model, and this was then tested on the remaining (X) epoch. This

process was repeated until all epochs had served as both test and training trials. We also repeated this procedure for each point in the epoch to produce a time-resolved orientation-selectivity function. To re-align the trials with the exact presented orientation, we reconstructed the stimulus representation[15] by multiplying the channel weights (9 channels × time × trial) against the basis set (180 orientations × 9 channels). This resulted in a 180 (-89° to 90°) Orientations × Trial × Time reconstruction. In order to average across trials, the orientation dimension was shifted so that 0º corresponded to the presented orientation in each trial.

These stimulus reconstructions were then averaged over conditions (Lag and Accuracy) and smoothed with a temporal Gaussian kernel with an 8 ms standard deviation[10,12]. To quantify orientation selectivity, we fit these responses with a Gaussian function with a constant offset using least square regression to quantify the amount of orientation selective activity, as per Equation 1, where A is the amplitude representing the amount of orientation selective activity, ϕ is the orientation on which the function is centred (in degrees), σ is the width (degrees), and C is a constant used to account for non-orientation selective baseline shifts.

Univariate orientation selectivity analysis

We used a univariate selectivity analysis[10] to determine the topography associated with orientation-selective activity for targets and distractors. Data were epoched in the same manner as in the forward encoding model where EEG activity was aligned with each stimulus presentation. We separated these epochs into target and distractor presentations to determine whether these stimulus types were processed

differently. All target presentations were used in training (1200 in total; 600 trials with two targets in each), together with a pseudo-random selection of the same number of distractor items. To determine the topography, we used a general linear model to estimate orientation selectivity for each sensor from the sine and cosine of the presentation orientation, and a constant regressor in each presentation. From the weights of the two orientation coefficients we calculated selectivity using Equation 7:

$$A = \sqrt{B_1 cos^2 + B_2 sin^2} \quad (7)$$

A was derived through permutation testing in which the design matrix was shuffled (N = 1000) and weights calculated. The non-permuted weights were ranked and compared with the permutation distribution, thus enabling calculation of the z-scored difference. To calculate group-level effects, cluster-based permutation testing (N=1500) across electrodes, time and participants, was implemented in Fieldtrip[57] to determine how the topographies differed between conditions.

Acknowledgements

This work was supported by the Australian Research Council (ARC) Centre of Excellence for Integrative Brain Function (ARC Centre Grant CE140100007) to JBM and EA, and by an ARC Discovery Project (DP170100908) to EA. JBM was supported by an ARC Australian Laureate Fellowship (FL110100103) and by the Canadian Institute for Advanced Research (CIFAR). JTE was supported by a Discovery Grant from the Natural Sciences and Engineering Research Council (Canada).

References

1. McLeod, P. Parallel processing and the psychological refractory period. *Acta Psychologica* **41,** 381–396 (1977).
2. Welford, A. T. The 'Psychological refractory period' and the timing of high-speed performance—A review and a theory. *British Journal of Psychology. General Section* **43,** 2–19 (1952).
3. Shapiro, K. L., Raymond, J. E. & Arnell, K. M. The attentional blink. *Trends Cogn Sci* **1,** 291–296 (1997).
4. Raymond, J. E., Shapiro, K. L. & Arnell, K. M. Temporary suppression of visual processing in an RSVP task: an attentional blink? . *Journal of Experimental Psychology: Human Perception and Performance* **18,** 849–860 (1992).
5. Broadbent, D. E. & Broadbent, M. H. P. From detection to identification: Response to multiple targets in rapid serial visual presentation. *Percept Psychophys* **42,** 105–113 (1987).
6. Williams, M. A., Visser, T. A. W., Cunnington, R. & Mattingley, J. B. Attenuation of Neural Responses in Primary Visual Cortex during the Attentional Blink. *J Neurosci* **28,** 9890–9894 (2008).
7. Marois, R., Yi, D.-J. & Chun, M. M. The Neural Fate of Consciously Perceived and Missed Events in the Attentional Blink. *Neuron* **41,** 465–472 (2004).
8. Vogel, E. K., Luck, S. J. & Shapiro, K. L. Electrophysiological evidence for a postperceptual locus of suppression during the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance* **24,** 1656–1674 (1998).
9. Garcia, J. O., Srinivasan, R. & Serences, J. T. Near-real-time feature-selective modulations in human cortex. *Current Biology* **23,** 515–522 (2013).
10. Myers, N. E. *et al.* Testing sensory evidence against mnemonic templates. *eLife Sciences* **4,** e09000 (2015).
11. Wolff, M. J., Jochim, J., Akyürek, E. G. & Stokes, M. G. Dynamic hidden states underlying working-memory-guided behavior. **20,** 864–871 (2017).
12. Tang, M. F., Smout, C. A., Arabzadeh, E. & Mattingley, J. B. Prediction error and repetition suppression have distinct effects on neural representations of visual information. *eLife Sciences* **7,** e33123 (2018).
13. Smout, C. A., Tang, M. F., Garrido, M. I. & Mattingley, J. B. Attention promotes the neural encoding of prediction errors. *PLoS Biol* **17,** e2006812 (2019).
14. Foster, J. J., Sutterer, D. W., Serences, J. T., Vogel, E. K. & Awh, E. The topography of alpha-band activity tracks the content of spatial working memory. *J Neurophysiol* **115,** 168–177 (2016).
15. Sprague, T. C. & Serences, J. T. Attention modulates spatial priority maps in the human occipital, parietal and frontal cortices. **16,** 1879–1887 (2013).
16. Brouwer, G. J. & Heeger, D. J. Decoding and reconstructing color from responses in human visual cortex. *J Neurosci* **29,** 13992–14003 (2009).
17. McAdams, C. J. & Maunsell, J. H. R. Effects of Attention on Orientation-Tuning Functions of Single Neurons in Macaque Cortical Area V4. *J Neurosci* **19,** 431–441 (1999).
18. Maunsell, J. Neuronal Mechanisms of Visual Attention. *Annu. Rev. Vis. Sci.* **1,** 373–391 (2015).
19. Ester, E. F., Sutterer, D. W., Serences, J. T. & Awh, E. Feature-Selective Attentional Modulations in Human Frontoparietal Cortex. *J Neurosci* **36,** 8188–8199 (2016).
20. Treue, S. & Trujillo, J. Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* (1999). doi:10.1038/21176
21. Maunsell, J. H. R. & Treue, S. Feature-based attention in visual cortex. *Trends Neurosci* **29,** 317–322 (2006).

22. Treue, S. Neural correlates of attention in primate visual cortex. *Trends Neurosci* **62,** 117 (2001).
23. Maloney, R. T., Jayakumar, J., Levichkina, E. V., Pigarev, I. N. & Vidyasagar, T. R. Information processing bottlenecks in macaque posterior parietal cortex: an attentional blink? *Exp Brain Res* **228,** 365–376 (2013).
24. Raymond, J. E., Shapiro, K. L. & Arnell, K. M. Similarity determines the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance* **21,** 653–662 (1995).
25. Duncan, J., Ward, R. & Shapiro, K. Direct measurement of attentional dwell time in human vision. *Nature* **369,** 313–315 (1994).
26. Chun, M. M. & Potter, M. C. A two-stage model for multiple target detection in rapid serial visual presentation. *Journal of Experimental Psychology: Human Perception and Performance* **21,** 109–127 (1995).
27. Di Lollo, V., Kawahara, J.-I., Shahab Ghorashi, S. M. & Enns, J. T. The attentional blink: resource depletion or temporary loss of control? *Psychol Res* **69,** 191–200 (2005).
28. Wyble, B., Bowman, H. & Nieuwenstein, M. The attentional blink provides episodic distinctiveness: Sparing at a cost. *Journal of Experimental Psychology: Human Perception and Performance* **35,** 787–807 (2009).
29. Akyürek, E. G. & Hommel, B. Target integration and the attentional blink. *Acta Psychologica* **119,** 305–314 (2005).
30. Akyürek, E. G. *et al.* Temporal target integration underlies performance at lag 1 in the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance* **38,** 1448–1464 (2012).
31. Asplund, C. L., Fougnie, D., Zughni, S., Martin, J. W. & Marois, R. The attentional blink reveals the probabilistic nature of discrete conscious perception. *Psychological Science* **25,** 824–831 (2014).
32. Dehaene, S. & Naccache, L. Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition* (2001). doi:10.1016/S0010-0277(00)00123-2
33. Nieuwenhuis, S. & de Kleijn, R. Consciousness of targets during the attentional blink: a gradual or all-or-none dimension? *Atten Percept Psychophys* **73,** 364–373 (2011).
34. Visser, T. A. W. Expectancy-Based Modulations of Lag-1 Sparing and Extended Sparing During the Attentional Blink. *Journal of Experimental Psychology: Human Perception and Performance* 1–17 (2015). doi:10.1037/a0038754
35. Alais, D., Leung, J. & Van der Burg, E. Linear Summation of Repulsive and Attractive Serial Dependencies: Orientation and Motion Dependencies Sum in Motion Perception. *J Neurosci* **37,** 4381–4390 (2017).
36. Fischer, J. & Whitney, D. Serial dependence in visual perception. *Nat Neurosci* **17,** 738–743 (2014).
37. Fritsche, M., Mostert, P. & de Lange, F. P. Opposite Effects of Recent History on Perception and Decision. *Current Biology* **27,** 590–595 (2017).
38. Dickinson, J. E., Almeida, R. A., Bell, J. & Badcock, D. R. Global shape aftereffects have a local substrate: A tilt aftereffect field. *J Vis* **10,** 5–5 (2010).
39. Dickinson, J. E., Morgan, S. K., Tang, M. F. & Badcock, D. R. Separate banks of information channels encode size and aspect ratio. *J Vis* **17,** 27 (2017).
40. Tang, M. F., Dickinson, J. E., Visser, T. A. W. & Badcock, D. R. The broad orientation dependence of the motion streak aftereffect reveals interactions between form and motion neurons. *J Vis* **15,** 4 (2015).
41. Visser, T. A. W., Bischof, W. F. & Di Lollo, V. Rapid serial visual distraction: task-irrelevant items can produce an attentional blink. *Percept Psychophys* **66,** 1418–1432 (2004).

42. Visser, T. A. W., Tang, M. F., Badcock, D. R. & Enns, J. T. Temporal cues and the attentional blink: a further examination of the role of expectancy in sequential object perception. *Atten Percept Psychophys* **76,** 2212–2220 (2014).
43. Visser, T. A. W. T1 difficulty and the attentional blink: Expectancy versus backward masking. *The Quarterly Journal of Experimental Psychology* **60,** 936–951 (2007).
44. Hanslmayr, S., Gross, J., Klimesch, W. & Shapiro, K. L. The role of alpha oscillations in temporal attention. *Brain Res Rev* **67,** 331–343 (2011).
45. Jun, J. J. *et al.* Fully integrated silicon probes for high-density recording of neural activity. *Nature* **551,** 232–236 (2017).
46. Pelli, D. G. The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spat Vis* **10,** 437–442 (1997).
47. Brainard, D. H. The Psychophysics Toolbox. *Spat Vis* **10,** 433–436 (1997).
48. Tang, M. F., Badcock, D. R. & Visser, T. A. W. Training and the attentional blink: limits overcome or expectations raised? *Psychonomic Bulletin & Review* **21,** 406–411 (2014).
49. Oostenveld, R. & Praamstra, P. The five percent electrode system for high-resolution EEG and ERP measurements. *Clinical Neurophysiology* **112,** 713–719 (2001).
50. Keil, A. *et al.* Committee report: Publication guidelines and recommendations for studies using electroencephalography and magnetoencephalography. *Psychophysiology* **51,** 1–21 (2014).
51. Bigdely-Shamlo, N., Mullen, T., Kothe, C., Su, K.-M. & Robbins, K. A. The PREP pipeline: standardized preprocessing for large-scale EEG analysis. *Front. Neuroinform.* **9,** B153 (2015).
52. Delorme, A. & Makeig, S. EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *J Neurosci Meth* **134,** 9–21 (2004).
53. Chaumon, M., Bishop, D. V. M. & Busch, N. A. A practical guide to the selection of independent components of the electroencephalogram for artifact correction. *J Neurosci Meth* **250,** 47–63 (2015).
54. Cheadle, S. *et al.* Adaptive Gain Control during Human Perceptual Choice. *Neuron* **81,** 1429–1441 (2014).
55. Kok, P., Mostert, P. & de Lange, F. P. Prior expectations induce prestimulus sensory templates. *Proc Natl Acad Sci USA* **26,** 201705652 (2017).
56. Blankertz, B., Lemm, S., Treder, M., Haufe, S. & Müller, K.-R. Single-trial analysis and classification of ERP components — A tutorial. *NeuroImage* **56,** 814–825 (2011).
57. Oostenveld, R., Martin, Maris, E. & Schoffelen, J.-M. FieldTrip: Open Source Software for Advanced Analysis of MEG, EEG, and Invasive Electrophysiological Data. *Computational Intelligence and Neuroscience* **2011,** 1–9 (2010).

Supplementary Figures

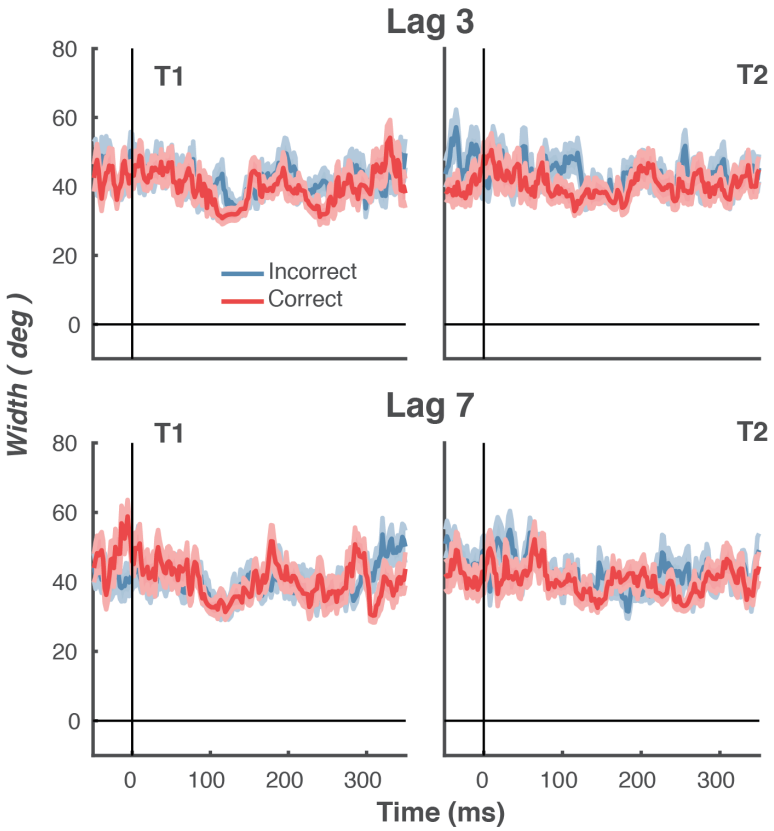

Supplementary Figure 1. Neural representations of orientation-selective information during the attentional blink in Experiment 2. **A.** Time course of measured fidelity of orientation selectivity for T1 and T2, given by the width (standard deviation) of the fitted Gaussian. Trials were scored as correct if the participant's response was within 30° of the presented orientation. Only trials in which participants accurately responded to T1 were included in the analysis. There were no significant differences between conditions (two-tailed cluster-permutation, alpha $p < .05$, cluster alpha $p < .05$, N permutations = 20,000).

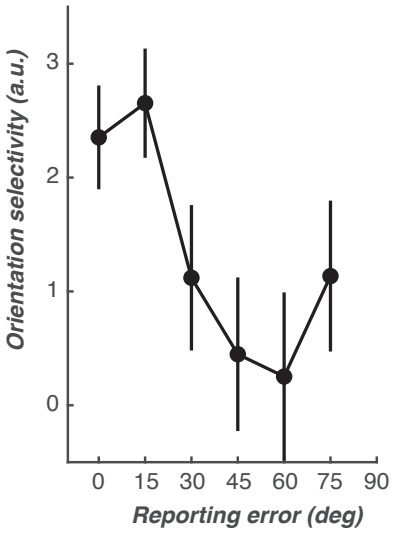

Supplementary Figure 2. Neural representations of orientation-selective information for T2 items presented at Lag 3 in Experiment 2, sorted by reporting error. Forward encoding results were averaged across early time points (100-150ms) in Figure 4A and were binned by the absolute difference between the presented and reported orientation (in 15° bins). These results were fitted with Gaussian functions to quantify across observers. Orientation selectivity was highest when participants reported the orientation within 30 degrees, and declined significantly with increasingly large reporting errors. Error bars indicate ∓1 standard error of the mean.