

Available online at www.sciencedirect.com

ScienceDirect

Journal homepage: www.elsevier.com/locate/cortex

Special issue: Research report

Hide and seek: Directing top-down attention is not sufficient for accelerating conscious access

Surya Gayet ^{a,*}, Iris Douw ^b, Vera van der Burg ^b,
Stefan Van der Stigchel ^b and Chris L.E. Paffen ^b

^a Donders Institute for Brain, Cognition and Behaviour, Nijmegen, The Netherlands

^b Experimental Psychology, Helmholtz Institute, Utrecht University, Utrecht, The Netherlands

ARTICLE INFO

Article history:

Received 14 February 2018

Reviewed 8 June 2018

Revised 20 July 2018

Accepted 27 August 2018

Published online 10 September 2018

Keywords:

Visual awareness

Consciousness

Attentional selection

Feature-based attention

Spatial attention

ABSTRACT

At any moment in time, we have a single conscious visual experience representing a minute part of our visual world. As such, the visual input stimulating our retinae is in continuous competition for reaching conscious access. Many complex cognitive operations can only be applied to consciously accessible visual information, thereby raising the question whether humans have the ability to select which parts of their visual input reaches consciousness. Top-down attention allows humans to flexibly assign more processing resources to certain parts of our visual input, making it a likely mechanism to volitionally bias conscious access. Here, we investigated whether directing top-down attention to a particular location or feature accelerates conscious access of an initially suppressed visual stimulus at the attended location, or of the attended feature.

We instructed participants to attend a spatial location (Experiment 1) or color (Experiment 2) for a speeded discrimination task, using a highly predictive cue. The predictive cues were highly effective in prompting sustained attention towards the cued location or color, as evidenced by faster discrimination of cued relative to uncued targets. We simultaneously measured detection times to interocularly suppressed probes that were either of the cued (i.e., attended) color/location or not, and were visually distinct from the targets used for the discrimination task. Despite our successful manipulation of top-down attention, suppressed probes were not released from suppression faster when they were presented at the attended location, or in the attended color. In contrast, when observers were cued to attend a color for locating targets of an ill-defined shape (inciting a broader attentional template), we did observe faster conscious access of probes in the attended color (Experiment 3). We discuss our findings in light of the specificity of attentional templates, and the inherent limitations that this poses for top-down attentional biases on conscious access.

© 2018 Elsevier Ltd. All rights reserved.

* Corresponding author. Montessorilaan 3, 6525 HR, Nijmegen, The Netherlands.

E-mail addresses: s.gayet@donders.ru.nl, surya.gayet@gmail.com (S. Gayet).

<https://doi.org/10.1016/j.cortex.2018.08.027>

0010-9452/© 2018 Elsevier Ltd. All rights reserved.

1. Introduction

Visual consciousness can be defined as a subjective experience in the visual modality, and is arguably the ultimate stage of visual processing. Although many differing theories of consciousness exist (e.g., Baars, 2005; Dehaene, Kerszberg, & Changeux, 1998; Dennett, 1991; Edelman and Tononi, 2008; Graziano, 2013), most theories share two commonalities: (1) consciousness is required for performing complex, demanding, and novel behavior, and (2) only a small part of our visual input eventually gives rise to a conscious experience (for a recent discussion on the extent to which high-level processes can be applied to non-conscious information, see Hassin, 2013; Hesselmann & Moors, 2015). This leads to the inevitable question: can an observer exert influence on selecting which part of the visual input gains access to consciousness? To address this question, we turn towards a well-known top-down selection mechanism: endogenous attention. Endogenous attention refers to an observer's ability to favor certain parts of the visual input (e.g., a location in space, or a particular color) at the expense of other parts (James, 1890). In other words, endogenous attention is used to delegate more cognitive resources to behaviorally relevant visual input at the expense of behaviorally irrelevant visual input (Desimone & Duncan, 1995). Here, we investigate whether visual input that is not accessible to consciousness, will gain accelerated access to consciousness if it is endogenously attended rather than unattended.

Before investigating the influence of endogenous attention on consciousness, it is important to consider that attention and consciousness are two highly intertwined phenomena. The degree to which attention and consciousness are dependent on each other has been a subject of much debate in the literature (for reviews, see Dehaene, Changeux, Naccache, Sackur, & Sergent, 2006; Koch & Tsuchiya, 2007; Lamme, 2003). Specifically, certain theorists have argued that attention (to a stimulus) is a necessary condition for consciousness (of that stimulus) to arise (e.g., Dehaene et al., 2006; Dennett, 1991; Graziano, 2013; Posner, 1994). Others have proposed that consciousness exists both in an attended and an unattended flavor. According to this view, we consciously experience a rich visual world (all of which constitutes our 'phenomenal consciousness'), but only the part of it that we attend to can be reported upon (which is referred to as 'access consciousness') (Block, 2011; Koch & Tsuchiya, 2007; Lamme, 2003). The relationship between consciousness and attention might even be a more intricate one, as evidence suggests that attention is required for conscious detection of certain types of information (e.g., detection and discrimination of novel stimuli) but not for others (e.g., gist perception) (Mack & Rock, 1998), and that attention selectively benefits conscious detection of stimuli that are relevant to the behavior of the observer (Rafal, Danziger, Grossi, Machado, & Ward, 2002). Arguably of most relevance to the current study, it has been proposed that non-conscious information can also be either attended or unattended (e.g., Koch, 2004; Koch & Tsuchiya, 2007, 2012; Lamme, 2003; Watanabe et al., 2011). In other words, the allocation of attentional resources towards stimuli that are not accessible to consciousness could affect the

extent to which they are processed. For instance, non-conscious stimuli require top-down attention in order to elicit a priming or cueing effect (e.g., Gayet, Van der Stigchel, & Paffen, 2014a; Kiefer & Brendel, 2006; Martens, Ansorge, & Kiefer, 2011; Naccache, Blandin, & Dehaene, 2002), but not for eliciting perceptual after-effects (Gilroy & Blake, 2005; Hofstoetter, Koch, & Kiper, 2004; Tsuchiya & Koch, 2005). Moreover, perceptual after-effects caused by non-conscious stimuli increase in magnitude when they are of an attended compared to an unattended color (Kanai, Tsuchiya, & Verstraten, 2006; Melcher, Papathomas, & Vidnyánszky, 2005). In contrast, such enhancement is not observed for non-conscious stimuli at a top-down attended compared to a top-down unattended location (Hsu, George, Wyart, & Tallon-Baudry, 2011). Taken together, it seems that despite the similarities between consciousness and attention, even the processing of non-conscious visual input can be modulated by top-down allocation of attentional resources. Here we investigated whether humans can capitalize on this by exerting volitional influence on what part of their visual world will gain access to consciousness.

In order to measure conscious access, we measured the time it took for initially perceptually suppressed visual input to overcome the suppression, so that it could be reported upon by observers. Perceptual suppression was achieved by means of continuous flash suppression (CFS): probe stimuli presented to one eye only, are interocularly suppressed by presenting a high-contrast dynamic pattern mask to the other eye (Tsuchiya & Koch, 2005). Measuring the time it takes for these probes to overcome suppression is referred to as the breaking continuous flash suppression (or b-CFS) paradigm, and allows for comparing the propensity of visual input in gaining access to consciousness, between different experimental manipulations (Jiang, Costello, & He, 2007; Stein, Hebart, & Sterzer, 2011; for a review, see Gayet, Van der Stigchel, & Paffen, 2014b). In the present study, we used this paradigm to compare the propensity of attended and unattended probes to overcome interocular suppression, so as to become accessible to consciousness.

Using the b-CFS paradigm renders the time at which a suppressed probe gains conscious access unpredictable, as suppression durations can vary by several seconds on a trial-to-trial basis (e.g., Stein et al., 2011). This has the advantage of being similar to real-world visual search, in which observers do not know beforehand at what point in time visual input will reach conscious access. Because of this temporal uncertainty, however, a manipulation of endogenous attention was required that incited participants to sustain their allocation of attentional resources throughout an entire trial. In this study, we tackled this issue by embedding a b-CFS paradigm within an endogenous cueing paradigm, creating a dual task: (1) In order to manipulate (and measure) the sustained orienting of endogenous attention, we required participants to perform a speeded discrimination task on a visible target, of which the location (Experiment 1) or color (Experiment 2) was predicted with high validity by a central cue. Crucially, this target could appear at any time during the 5 sec trial, inciting participants to sustain their (spatial or feature-based) attention throughout the entire trial. (2) In order to measure the influence of sustained endogenous attention on conscious access,

participants were simultaneously required to perform a speeded detection task on an initially interocularly suppressed probe that was presented on a minority of trials, either at the attended location (or in the attended color) or not. This suppressed probe was presented at the start of the trial, but could escape from suppression (and become reportable) at any time during the 5 sec trial. The current approach thus maximizes the potential impact of endogenous attention on suppressed visual input, as it allows endogenous attention to be sustained during the entire time that the probe is suppressed and thus not (yet) detected.

When observers deploy endogenous attention to favor relevant visual input over irrelevant visual input, it is assumed that an attentional ‘template’ is kept available to the visual system using working memory. Specifically, spatial attention (i.e., looking for something at a particular location in space) is assumed to require maintenance of a location in spatial working memory (e.g., LaBar, Gitelman, Parrish, & Mesulam, 1999; for a review, see Awh & Jonides, 2001). Indeed, both spatial attention (Buracas & Boynton, 2007; Gandhi, Heeger, & Boynton, 1999) and spatial working memory (Merrikhi et al., 2017) enhance the neural response to visual input presented at that location. Similarly, feature-based attention (e.g., looking for something red) is assumed to require maintenance of a feature in visual working memory (Bundesen, Habekost, & Kyllingsbæk, 2005; de Fockert, Rees, Frith, & Lavie, 2001; Gunseli, Meeter, & Olivers, 2014; Wolfe, 1994). Accordingly, both feature-based attentional templates (Treue & Trujillo, 1999; for a review, see Treue, 2001) and visual working memory (Gayet et al., 2017) enhance the neural response to matching visual input. A further similarity is that both attention (Greenberg, Esterman, Wilson, Serences, & Yantis, 2010) and working memory (Postle, Stern, Rosen, & Corkin, 2000) have modality-independent frontal control regions, in addition to more posterior content-specific regions that are associated with sensory processing. Taken together, it can be inferred that instructing participants to attend to a particular location or color, causes participants to maintain this particular location or color in working memory. Using a b-CFS paradigm, we recently demonstrated that a color that is actively maintained in visual working memory accelerates conscious access of concurrently presented visual input of the same color category (Gayet, Paffen, & Van der Stigchel, 2013; Gayet, van Maanen, Heilbron, Paffen, & Van der Stigchel, 2016; van Moorselaar et al., 2018). Considering that feature-based attention draws upon visual working memory, we hypothesized that attending to a particular color would accelerate conscious access of initially suppressed visual input that matches this color (Experiment 2). Extrapolating this prediction to the spatial domain, led us to hypothesize that attending to a particular location, which relies on maintaining said location in spatial working memory, would accelerate conscious access of initially suppressed visual input at that location (Experiment 1).

Because we were investigating whether (alternative hypothesis) or not (null hypothesis) sustained endogenous attention would impact conscious access, we used a Bayesian approach, which allows for finding statistical support for the alternative hypothesis, as well as for the null hypothesis.

2. Experiment 1: methods

2.1. Participants and stopping rule

Participants ($N = 28$) were recruited at Utrecht University, and rewarded for participation with course credits or monetary compensation. Eventually, 25 participants (7 males) with a mean age of 23.6 years ($SD = 3.3$) were included in the analyses (see Section 4.1 for further elaboration on participant exclusion). All participants had (corrected to) normal vision, including stereopsis (tested with the TNO test for stereoscopic vision, 12th edition; Walraven, 1972) and color vision (tested with the Ishihara color blindness test plates; Ishihara, 1960). Participants provided written informed consent before participation. This study was approved by our faculty's local ethics committee as part of an overarching ethical application covering non-invasive behavioral experiments with healthy participants.

We set out to test a minimum of 20 participants, and collect participants until a Bayes factor (BF) of 6 was obtained for either the null hypothesis ($BF_{01} \geq 6$, supporting that attention has no effect on response times to suppressed targets) or the alternative hypothesis ($BF_{10} \geq 6$, supporting that attention does have an effect on response times to suppressed targets; non-directional), or until the pattern of results became robust to the addition of new participants (in any of those two directions). Eventually, a BF above 3 for either the null or the alternative hypothesis was regarded as substantial evidence in favor of that hypothesis (Dienes, 2011; Jeffrey, 1961; Lee & Wagenmakers, 2014). Please note that, for the sake of consistency, all BFs are reported so as to reflect evidence for the alternative hypothesis (i.e., BFs below $\frac{1}{3}$ for the alternative hypothesis should be interpreted as BFs above 3 for the null hypothesis).

2.2. Procedure and design

Participants took part in three practice sessions (24 trials each) before participating in the actual experiment (320 trials, separated into 8 blocks of 40 trials each). For the first practice task, participants were instructed to fixate the central fixation dot, and press the spacebar as fast as possible as soon as a red circle (the probe) would become visible. This probe was presented between .5 and 1 sec (variable stimulus-onset asynchrony) after onset of the CFS masks and was thus initially interocularly suppressed. The probe was ramped up from zero to full intensity over the course of 1 sec (in line with standard practice in the b-CFS paradigm), minimizing stimulus-driven reversals of ocular dominance at probe onset, and maximizing manipulation-driven reversals of ocular dominance throughout the trial duration. After participants provided a response, or after a 5-sec response deadline had elapsed since probe onset, the initially suppressed probe was removed from the screen, but the CFS sustained until 5.5 sec had elapsed. In the second practice task participants were instructed to report the orientation (upright or inverted) of a T-shaped target, that could appear left or right of fixation, as fast as possible. Participants were instructed that, preceding this T-target, an arrow cue would help them in reporting the orientation of the

T-target as fast as possible, as it would correctly predict its location (left or right hemi-field) in 87.5% of the trials (i.e., 224 valid trials). Accordingly, in the remaining 12.5% of the trials the arrow would incorrectly predict the location of the subsequent targets (i.e., 32 invalid trials). The arrow cue was presented for 800 msec, 800 msec prior to CFS onset. Crucially, the T-target could appear at any time between .5 and 4.5 sec after CFS onset, inciting participants to maintain an active search strategy (i.e., attending to the cued hemi-field) throughout the entire duration of a trial. The third practice task (which was identical to the main experiment) consisted of a combination of the first two tasks: participants were instructed that on each trial, they should (1) utilize the arrow cue to help them report the orientation of a T-target as fast as possible, and (2) press spacebar as soon as the red circular probe would appear. The red probe appeared in 20% of the trials (64 out of 320).

In sum, the participants were – from their point of view – performing a double task: (1) attending to a hemi-field to report the orientation of a T-target, and (2) reporting the incidental appearance of a probe. From the experimenters' point of view, however, participants were performing one of two possible trial types, which were randomly intermixed within experimental blocks: In the first trial type (right panel of Fig. 1), the arrow cue was followed by a T-target during CFS

presentation, while no red probe was presented. These 256 trials were used to manipulate participants' spatial attention towards the cued hemi-field, and (in the analysis stage) also allowed us to establish empirically whether this manipulation was successful. In the second trial type (left panel of Fig. 1), the arrow cue was followed by a red probe appearing shortly after CFS onset. These 64 trials were used to measure the influence of spatial attention (evoked by the arrow cue) on suppression durations of the initially suppressed probes. This subdivision in two trial types allowed us to investigate the participants' responses to the two target types (i.e., suppressed probes, and T-targets), with no interference from the other target type. In order to maintain the participants' impression of performing a double task, a T-target was presented after a response had been provided to the red probe, or after 5 sec had elapsed and no response to the red probe had been provided. The position of these 'sham' T-targets was randomized (but not counter-balanced), and responses to these 'sham' T-targets were not considered for further analyses.

All experimental factors were manipulated within participants and were equated between experimental blocks. Considering the first trial type (256 in total), which was designed to manipulate participants' spatial attention, the predictive arrow cue could point either left (50%) or right (50%), the subsequent T-target could appear either at the cued

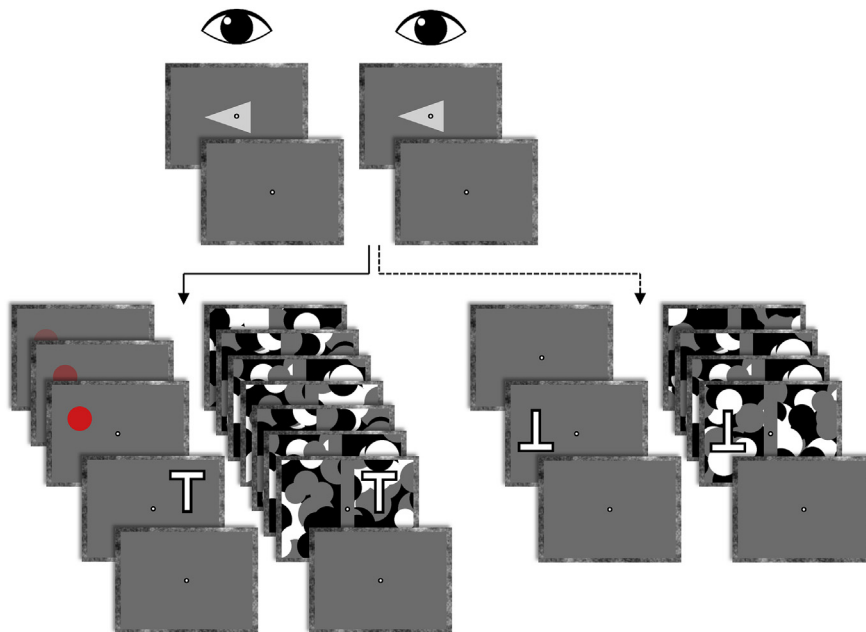


Fig. 1 – A schematic depiction of two trial types in Experiment 1. Each trial started with an arrow cue pointing left or right, which was highly predictive of the subsequent T-target location. Next, a dynamic pattern was presented to one eye, and a dichoptically presented T-target could appear at any moment within the next 5 sec. Participants were instructed to report the orientation (upright or inverted) of the T-target as fast and accurately as possible. Additionally, participants were instructed that in some trials a red circle could appear (left timeline), in which case they should press the spacebar as fast as possible. The left timeline shows a trial with a red circle (the initially suppressed probe) in the cued hemi-field, followed by an upright T-target in the uncued hemi-field. Trials including a red probe were used to assess whether directing top-down attention to a location affected detection of suppressed visual input at that location. The response times to T-targets that followed the red probe were omitted from further analyses. The right timeline shows a trial with an inverted T-target on the cued location (i.e., a trial with a valid cue). These trials were used to assess whether participants used the cue to speed up their responses to the T-target, by comparing trials in which the cue was valid (i.e., the T-target was presented to the cued hemi-field) with trials in which the cue was invalid (i.e., the T-target was presented to the uncued hemi-field).

location (87.5%, or 224 valid trials) or the uncued location (12.5%, or 32 invalid trials) and was presented either upright (50%) or inverted (50%). The orientation of the T-target and the pointing direction of the arrow cue were counterbalanced within valid trials and invalid trials. The moment in time at which the T-target would appear was chosen from a linear division of 32 time-points between .5 and 4.5 sec, and the exact angular location of the T-target was chosen from a linear division of 8 angular locations within 30° of the horizontal midline. These two factors were not counterbalanced, but the prevalence of factor-levels was equated within blocks. Considering the second trial type (64 in total), which was designed to measure the influence of spatial attention on suppression durations, the arrow cue could point either left (50%) or right (50%), the initially suppressed probe could appear either at the cued location (50%, or 32 attended trials) or the uncued location (50%, or 32 unattended trials) and was either presented to the left eye (50%) or the right eye (50%) of the participant. The prevalence of these factor-levels was fully counterbalanced, and the arrow was therefore not predictive of the subsequent location of the probe. The exact moment in time at which the initially suppressed probe appeared was randomly chosen from a linear division of 8 time-points between .5 and 1.0 sec, and the exact angular location of the suppressed probe was chosen from a linear division of 8 angular locations within 30° of the horizontal midline. These two factors were not counterbalanced, but the prevalence of factor-levels was equated within blocks. The experiment was subdivided into eight blocks (separated by mandatory breaks), each comprising 32 trials of the first type and 8 trials of the second type (in random order), with all factors counterbalanced within block as described above. As a final note, it follows from the prevalence of the two trial types that after a certain cue (say, a leftwards pointing arrow) there was a combined 80% chance that any of the to-be-reported stimuli (i.e., suppressed probe or T-target) would be presented at the cued location.

2.3. Stimuli and apparatus

Participants viewed a linearized CRT monitor (22" LaCie Electron Blue IV, 1024 × 768 pixels) through a mirror stereoscope, which ensured separate stimulation of the two eyes. A chinrest was used to maintain the effective viewing distance at 57 cm. Responses were collected using an Apple keyboard, and the temporal resolution of response time registration was limited by the refresh rate of the monitor (100 Hz). MATLAB (Release R2007b; The Mathworks, Natick, MA, USA) and Psychophysics Toolbox (Brainard, 1997; Pelli, 1997) were used for stimulus presentation. In order to promote binocular fusion of the complementary images, identical Brownian noise frames (width of 3.7 degrees of visual angle – dva) were presented to both eyes, delimiting the gray (16 cd/m²) presentation area on which all stimuli were presented. The presentation area consisted of a rectangular region with sides of 10 dva. At all times, a black (<.1 cd/m²; .04 dva) and white (32 cd/m²; .2 dva) fixation bullseye was presented at the center of the presentation areas.

The arrow cue was a gray (24 cd/m², 50% Weber luminance contrast with the background) isosceles triangle pointing left

or right, with a height of 1.25 dva, and a width of 3.13 dva. The triangle's center of mass was centered on fixation.

The masks used for eliciting CFS were composed of overlapping white, gray and black circles of different sizes (diameters ranging between .8 and 1.7 dva) and were chosen such as to maximally suppress the circular probe. These masks were similar to those used in earlier studies (e.g., Gayet, Paffen, Belopolsky, Theeuwes, & Van der Stigchel, 2016; Stein et al., 2011). On each block, 20 new masks were generated. During a trial, these masks were presented for 100 msec each (10 Hz) in random order, with the restriction that the same mask was never presented twice in succession. The masks were divided in two separate maskeable hemi-fields by a vertical saving of 2 dva. Because reversals of ocular dominance tend to follow object continuations (Kaufman, 1963), this saving increased the independency of suppression between hemi-fields (Maruya & Blake, 2009; for a similar approach, see Gayet, van Maanen, et al., 2016).

The suppressed probe consisted of a red circle with a diameter of 1.25 dva, and a Weber luminance contrast of 84% with the background (CIE values, as measured with a Photo Research Inc. PR-650 SpectraScan colorimeter/telephotometer: $x = .635$, $y = .343$, $z = 2.61$ cd/m²). The probe was presented monocularly, contra-ocular to the masks. The T-Target's consisted of a black outer-T with a length and width of 1.35 dva, and an edge width of .13 dva encapsulating a white inner-T. Using a black and white T-target was necessary to ensure its visibility irrespective of whether it happened to appear on top of a white or a black part of the masks. The T-Target was presented binocularly. The probes and T-targets were presented at an eccentricity of 3.5 dva from fixation, and could appear left or right of fixation, within 30 angular degrees from the horizontal midline.

3. Experiment 1: results

3.1. Participant exclusion and data preparation

Three participants failed to report the presence of a suppressed probe within the 5 sec trial duration on more than 50% of the trials, in one or more conditions (total hit-rates of 36%, 25%, and 39%). Because missing more than 50% of response time data within a condition precludes computation of a median response time for that condition, these participants were excluded from further analyses (for a similar approach, see Gayet, Paffen, et al., 2016). The remaining 25 participants had an average hit-rate of 83% (SD = 15) for detecting the suppressed probe. Trials in which the suppressed probe was not detected were regarded as yielding infinitely long response times (RTs) and were included in the computation of the median RT (thereby retaining the counterbalancing of all experimental conditions). For the sake of consistency, median response times were also used to summarize participants' response times to T-targets.

The computation of the median response time to probes in the attended condition was based on 32.0 trials (SD = .2), of which 26.8 (SD = 4.5) were hits and therefore yielded an actual RT value. In the unattended condition, the computation of the median response times to probes was based on 32.0 trials

(SD = .0), of which 26.4 (SD = 5.0) were hits and therefore yielded an actual RT value. A total of 1 trial (over all participants) was excluded from these analyses because it was responded to within 250 msec after probe onset, which is too fast to reflect a genuine response to the probe.

Kolmogorov–Smirnov tests indicated that the response time data violated the assumption of normality in all conditions (all p 's < .05). Because latency-normalized response time differences (see Gayet & Stein, 2017) still violated the assumption of normality in one condition, we chose to apply a \log_e -transformation to all raw response times (i.e., before computing the median response time), after which none of the conditions violated the assumption of normality (all p 's > .2) (Bush, Hess, & Wolford, 1993). This allowed us to conduct Bayesian t-tests that assume normally distributed data. The \log_e -transformed data is only used for statistical testing; to facilitate interpretation, however, all response times depicted in graphs and reported in text reflect the raw response times (in seconds or milliseconds).

For the analyses of response times to the T-targets, only correct responses were included. Participants were 96.4% (SD = 2.0) correct on valid trials, and 97.3% (SD = 3.2) correct on invalid trials. According to a Bayesian t-test, the observed data it is about twice as likely to reflect similar than it is to reflect dissimilar accuracy scores between conditions, but this is inconclusive evidence, $BF_{10} = .48$.

3.2. Manipulation of attention

First, we aimed to assess whether the predictive cues were successful in directing participants' attention to the cued hemi-field (Fig. 2A). Participants were 23 msec (SD = 22) faster at reporting the orientation of a T-target following a valid cue ($M = 522$ msec, $SD = 38$) than following an invalid cue ($M = 546$ msec, $SD = 42$). Numerically, 22 out of 25 participants showed faster response times after valid trials than after invalid trials (the solid lines in the left panel of Fig. 2A). A Bayesian paired-samples t-test confirmed that this difference was extremely reliable, $BF_{10} = 1747$, and robust to the width of the prior and the addition of participants (Fig. 2A, right panel).

Fig. 5A illustrates the stability of the cueing effect throughout the 5 sec trial. Although the cueing effect was numerically larger when the T-target appeared on the first half of the 5 sec trial ($M = 30$ msec, $SD = 27$) than in the second half ($M = 20$ msec, $SD = 28$), the evidence supporting this difference is inconclusive, and is more than twice as likely to reflect similar than it is to reflect dissimilar cueing effects between the first and second half of a trial, $BF_{10} = .61$. Similarly, Fig. 5B illustrates the stability of the cueing effect over the course of the experiment. While the cueing effect was numerically larger in the second half of the experiment ($M = 32$ msec, $SD = 29$) than in the first half ($M = 26$ msec, $SD = 24$), the evidence supporting this difference is inconclusive, and is almost twice as likely to reflect similar than it is to reflect dissimilar cueing effects between experimental halves, $BF_{10} = .61$.

Taken together, these data show that the predictive cues were successful in manipulating participants' attention, and consistently directed their spatial attention to the cued hemi-field, throughout the entire trial and over the course of the entire experiment.

3.3. Influence of attention on suppression durations

Now that the manipulation of attention through predictive central cueing is established to be successful and robust across time, we can distinguish between suppressed probes that were presented to either the attended (i.e., cued) or the unattended (uncued) hemi-field (Fig. 2B). Unlike the response times to the dichoptically presented T-targets, however, detection times to suppressed probes did not reliably differ depending on whether they were presented to the attended ($M = 1890$ msec, $SD = 778$), or the unattended hemi-field ($M = 1911$ msec, $SD = 888$). A Bayesian t-test confirmed that the observed data were about 4.5 times as likely to support the null hypothesis (no difference between conditions) than it was to support the alternative hypothesis (a difference between conditions), $BF_{10} = .23$. The evolution of the BF after addition of each participant (Fig. 2B, right panel) shows that this null effect is robust to the addition of further participants and is even more pronounced when a less informative prior is chosen.

Possibly, not all participants were equally susceptible to the attention manipulation (see the slopes in Fig. 2A, left panel). As such, we hypothesized that if the attention manipulation had any effect on the detection of suppressed probes, it should be most pronounced for participants that exhibited the strongest effect of the cue on response times to the T-target. To investigate this possibility, we performed a Bayesian pairwise correlation analyses between the response time difference on T-targets (valid vs invalid, reflecting the success of the attention manipulation) and the response time difference to the suppressed probes (attended vs unattended; Fig. 2C). This analysis revealed that the success of the attention manipulation was not related to the influence of spatial attention on the detection of suppressed probes; the observed data were about 4 times more likely to support the null hypothesis (no correlation) than to support the existence of this correlation, $R = .04$, $BF_{10} = .25$.

Finally, the hit-rates for attended ($M = 83.8\%$, $SD = 14.3$) and unattended probes ($M = 82.5\%$, $SD = 15.6$) did not differ ($BF_{10} = .28$, or 3.6 times more evidence for the null than for the alternative hypothesis), showing that sustained attention did not impact the hit-rate to suppressed probes. Taken together, these data demonstrate that orienting attention to a location in space is not sufficient for enhancing conscious access of suppressed visual input at that location.

4. Experiment 2: methods

4.1. Rationale

The goal of Experiment 2 was to investigate whether sustained endogenous attention to a feature (i.e., color) accelerates conscious access of (initially) suppressed visual input of the attended color, compared to visual input of the unattended color. Kanai et al. (2006) showed that interocularly suppressed gratings elicited more pronounced after-effects when they were composed of an attended feature compared to an unattended feature, whereas spatial attention elicited no such differential effect. Here, we address the question whether such attentional enhancement of interocularly

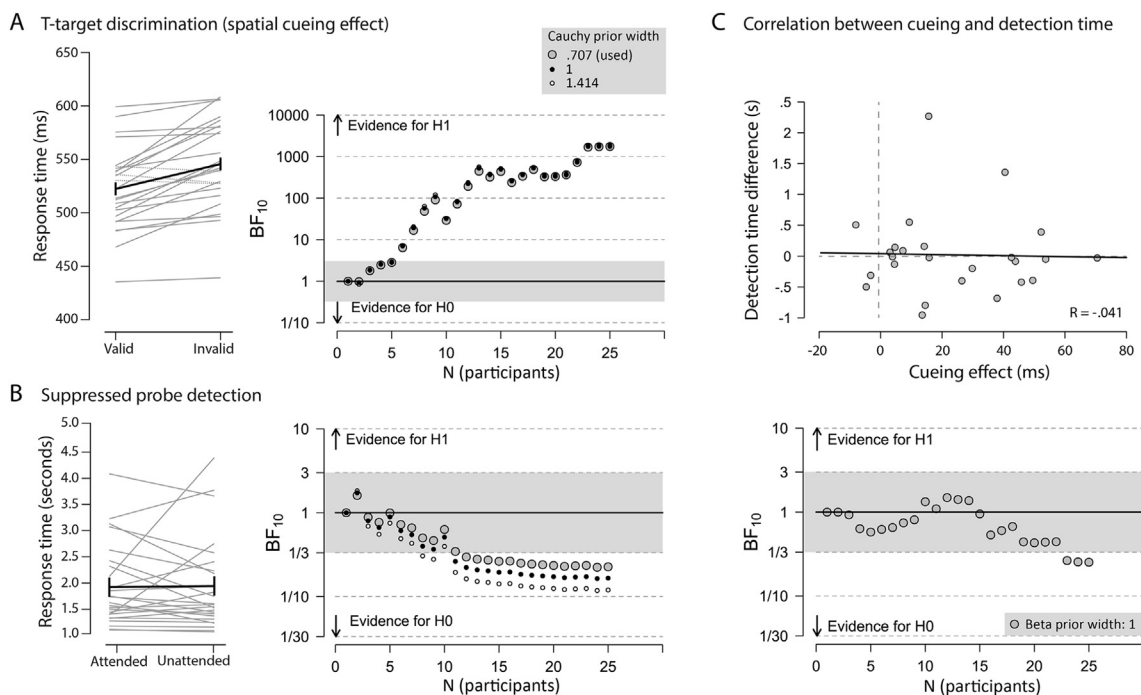


Fig. 2 – Results of Experiment 1: no influence of spatial attention on conscious access. Panel A relates to the speeded discrimination (upright or inverted) of a dichoptically presented T-target, depending on whether it was presented at the cued hemi-field (i.e., valid) or not (i.e., invalid). The left part shows response times in milliseconds, with individual participants in gray, and the group average in black. The right part shows the evolution of the BF for the difference in response times between trials in which the cue was valid versus invalid (i.e., the cueing effect). Panel B, is coded similarly to Panel A, but relates to the detection times of initially interocularly suppressed probes (in seconds), which were either presented at the attended (i.e., cued) or the unattended (i.e., uncued) hemi-field. Panel C shows the relation (or the absence thereof) between the influence of the cue on T-target discrimination (x-axis), and on detection of the initially interocularly suppressed probe (y-axis). Error bars in the response time graphs depict 95% confidence intervals after removal of the between-subject variance (Cousineau, 2005; using the correction proposed by Morey, 2008). Shaded areas in the sequential analysis graphs correspond to BFs that convey insufficient evidence to support either the null or the alternative hypothesis (Dienes, 2011; Jeffreys, 1961; Lee & Wagenmakers, 2014). Statistical analyses were conducted on \log_e -transformed response times (Bush et al., 1993; Gayet & Stein, 2017). Sequential analysis plots were adapted from JASP 0.8.2.0 output (JASP Team, 2017).

suppressed visual input can accelerate conscious access to the suppressed visual input. If this is the case, we should observe faster response times to initially interocularly suppressed probes when they are of the attended compared to the unattended color. In order to promote comparisons with Experiment 1, which measured the influence of spatial attention on conscious access, we kept the experimental paradigm as similar as possible to that of Experiment 1. Below, we describe only the modifications that were made for Experiment 2.

4.2. Participants

Thirty new participants were recruited at Utrecht University, and rewarded for participation with course credits or monetary compensation. Eventually, 27 participants (six males) with a mean age of 23.5 years ($SD = 2.8$) were included in the analyses (see Section 5.1 for further elaboration on participant exclusion). All participants had (corrected to) normal vision, including stereopsis (tested with the TNO test for stereoscopic vision, 12th edition; Walraven, 1972) and color vision (tested with the Ishihara color blindness test plates; Ishihara, 1960).

4.3. Additional methods, stimuli, and procedure

The first trial type (256 trials) was designed to manipulate (and assess) participants' feature-based attention towards either the color red or the color blue (Fig. 3, right panel). Participants were instructed to report the location (left or right hemi-field) of an inverted T-target as fast as possible, while a distractor (an upright T-target) was presented to the opposite hemi-field. The 87.5% valid cue that preceded the T-target consisted of a verbal label (i.e., the word “red” or “blue”), that predicted the color of the inverted T-target. The additional distractor was added to increase competition in the visual field, such as to increase the usefulness of deploying feature-based attention to perform the speeded task. Taken together, the predictive cue could be either “red” (50%) or “blue” (50%), the subsequent inverted T-target was either of the cued color (87.5%, or 224 valid trials) or the uncued color (12.5%, or 32 invalid trials) and was presented either to the left (50%) or the right hemi-field (50%). The color and location of the inverted T-target were counterbalanced within valid trials and invalid trials. The color of the distractor (upright T-target) was always of a

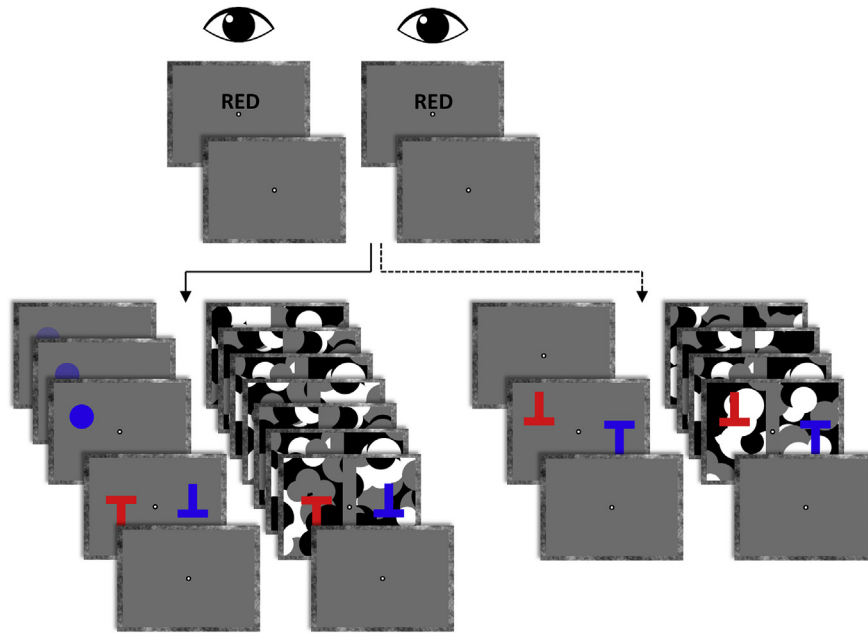


Fig. 3 – A schematic depiction of two trial types in Experiment 2. Each trial started with the word “red” or “blue”, which was highly predictive of the subsequent T-target color. Next, a dynamic pattern was presented to one eye, and two dichoptically presented T-targets could appear at any moment within the next 5 sec. Crucially, the T-targets appearing left and right were always of a different color (red or blue) and orientation (upright or inverted). Participants were instructed to report the location (left or right hemi-field) of the inverted T-target as fast and accurately as possible. Additionally, participants were instructed that in some trials a colored circle could appear (left timeline), in which case they should press spacebar as fast as possible. The left timeline shows a trial with a colored circle (the initially suppressed probe) presented in the uncued color, followed by an inverted T-target of the uncued color as well. Trials including a suppressed probe were used to assess whether directing top-down attention to a color affected detection of suppressed input of that color. The response times to T-targets that followed the suppressed probe were omitted from further analyses. The right timeline shows a trial with an inverted T-target presented in the cued color (i.e., a trial with a valid cue). These trials were used to assess whether participants were using the cue to speed up their response to the inverted T-target, by comparing trials in which the cue was valid (i.e., the inverted T-target was presented in the cued color) with trials in which the cue was invalid (i.e., the T-target was presented in the uncued color).

different color than the inverted T-target. The exact angular location of the distractor was chosen from a linear division of 8 angular locations within 30° of the horizontal midline.

The second trial type (64 trials) was designed to measure the influence of feature-based attention on suppression durations (Fig. 3, left panel). As in Experiment 1, participants were required to provide a speeded response upon probe detection, but now the probes could be either of the cued (i.e., attended) color, or of the uncued (i.e., unattended) color. Probes of the attended or unattended color were presented with equal prevalence (50%, or 32 trials each).

The colors of the red-colored targets ($x = .635$, $y = .343$, $z = 2.61$ cd/m², 84% Michelson contrast with the background), and the blue-colored targets ($x = .147$, $y = .068$, $z = 2.64$ cd/m², 84% Michelson contrast with the background), were chosen such as to be perceptually equiluminant based on the group data of one of our recent previous studies (Gayet, Paffen, et al., 2016).

5. Experiment 2: results

5.1. Participants and data exclusion

Two participants were excluded from further analyses because they failed to report the presence of a suppressed probe within the 5 sec trial duration on more than 50% of the trials, in one or more conditions (total hit-rates of 39% and 31%). The remaining 27 participants had an average hit-rate of 77% (SD = 15) for detecting the suppressed probe. One additional participant was excluded from further analyses because his/her false alarm rate for reporting the suppressed probe (16 times when there was none) was more than 4 standard deviations above the group mean.

The computation of the median response time to probes in the attended condition was based on 32.0 trials (SD = .0), of which 24.2 (SD = 5.5) were hits and therefore yielded an

actual RT value. In the unattended condition, the computation of the median response times to probes was based on 31.8 trials ($SD = .5$), of which 24.0 ($SD = 6.1$) were hits and therefore yielded an actual RT value. A total of six trials (over all participants) were excluded from these analyses because they were responded to within 250 msec after probe onset, which is too fast to reflect a genuine response to the probe.

As in Experiment 1, Kolmogorov–Smirnov tests indicated that the response time data violated the assumption of normality in all conditions (all p 's < .05). After \log_e -transformation of raw response times the assumption of normality was no longer violated (all p 's > .1). Again, the \log_e -transformed data is only used for statistical testing; to facilitate interpretation, all response times depicted in graphs and reported in text reflect the raw response times (in seconds or milliseconds).

For the analyses of response times to the T-targets, only correct responses were included. Participants were 96.0% ($SD = 2.9$) correct on valid trials, and were reliably less accurate on invalid trials 83.3% ($SD = 14.2$), $BF_{10} = 857$.

5.2. Manipulation of attention

First, we aimed to assess whether the predictive cues were successful in directing participants' attention to the cued hemi-field (Fig. 4A). Participants were 65 msec ($SD = 52$) faster at reporting the location of the inverted T-target following a valid cue ($M = 613$ msec, $SD = 62$) than following an invalid cue ($M = 679$ msec, $SD = 73$). Numerically, 25 out of 27 participants showed faster response times in valid trials than in invalid trials (the solid lines in the left panel of Fig. 4A). A Bayesian paired-samples t -test confirmed that this difference was extremely reliable, $BF_{10} = 44301$, and robust to the width of the prior and the addition of participants (Fig. 4A, right panel).

Fig. 5C illustrates the stability of the cueing effect throughout the 5 sec trial. The cueing effect was numerically larger when the T-targets appeared in the second half of the trial ($M = 78$ msec, $SD = 77$) than in the first half ($M = 63$ msec, $SD = 604$), but this difference was not reliable, as a Bayesian t -test provides more than 3 times more evidence for a null effect than for a difference in cueing effect between experimental halves, $BF_{10} = .30$. Fig. 5D illustrates the stability of the cueing

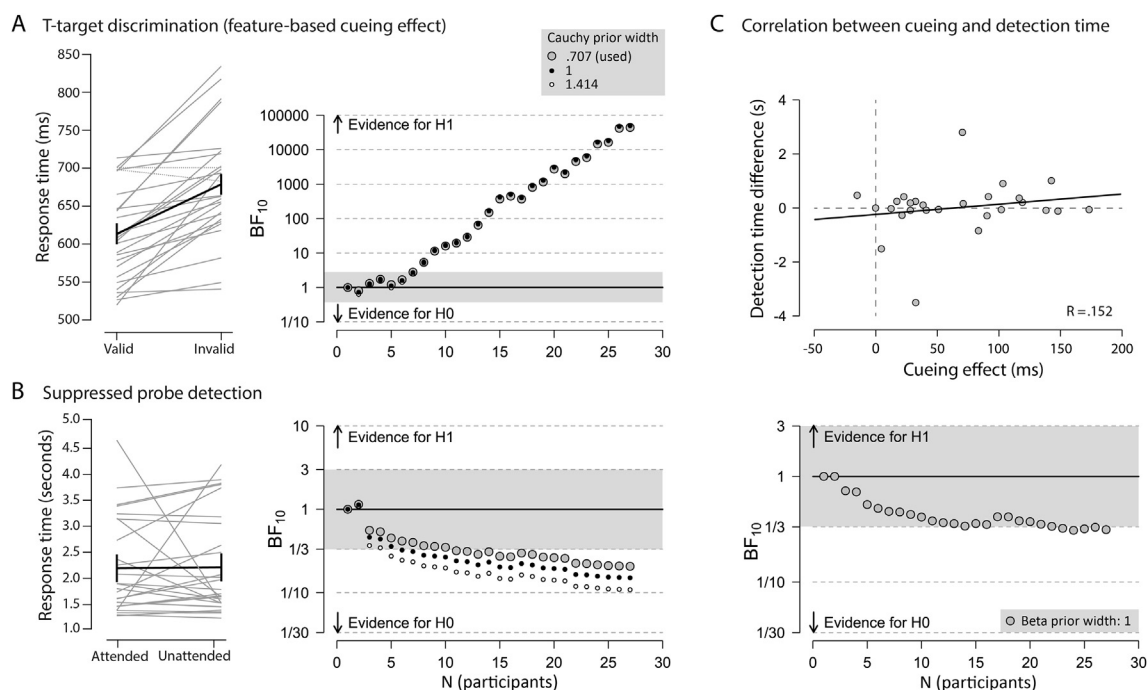


Fig. 4 – Results of Experiment 2: no influence of feature-based attention on conscious access. Panel A relates to the speeded discrimination (i.e., presented to the left or right hemi-field) of a dichoptically presented inverted T-target, depending on whether it was of the cued color (i.e., valid) or not (i.e., invalid). The left part shows response times in milliseconds, with individual participants in gray, and the group average in black. The right part shows the evolution of the BF for the difference in response times between trials in which the cue was valid versus invalid (i.e., the cueing effect). Panel B, is coded similarly to Panel A, but relates to the detection times of initially interocularly suppressed probes (in seconds), which were either of the attended (i.e., cued) or the unattended (i.e., uncued) color. Panel C shows the relation (or the absence thereof) between the influence of the cue on T-target discrimination, and on detection of the initially interocularly suppressed probe. Error bars in the response time graphs depict 95% confidence intervals after removal of the between-subject variance (Cousineau, 2005; using the correction proposed by Morey, 2008). Shaded areas in the sequential analysis graphs correspond to BFs that convey insufficient evidence to support either the null or the alternative hypothesis (Dienes, 2011; Jeffreys, 1961; Lee & Wagenmakers, 2014). Statistical analyses were conducted on \log_e -transformed response times (Bush et al., 1993; Gayet & Stein, 2017). Sequential analysis plots were adapted from JASP 0.8.2.0 output (JASP Team, 2017).

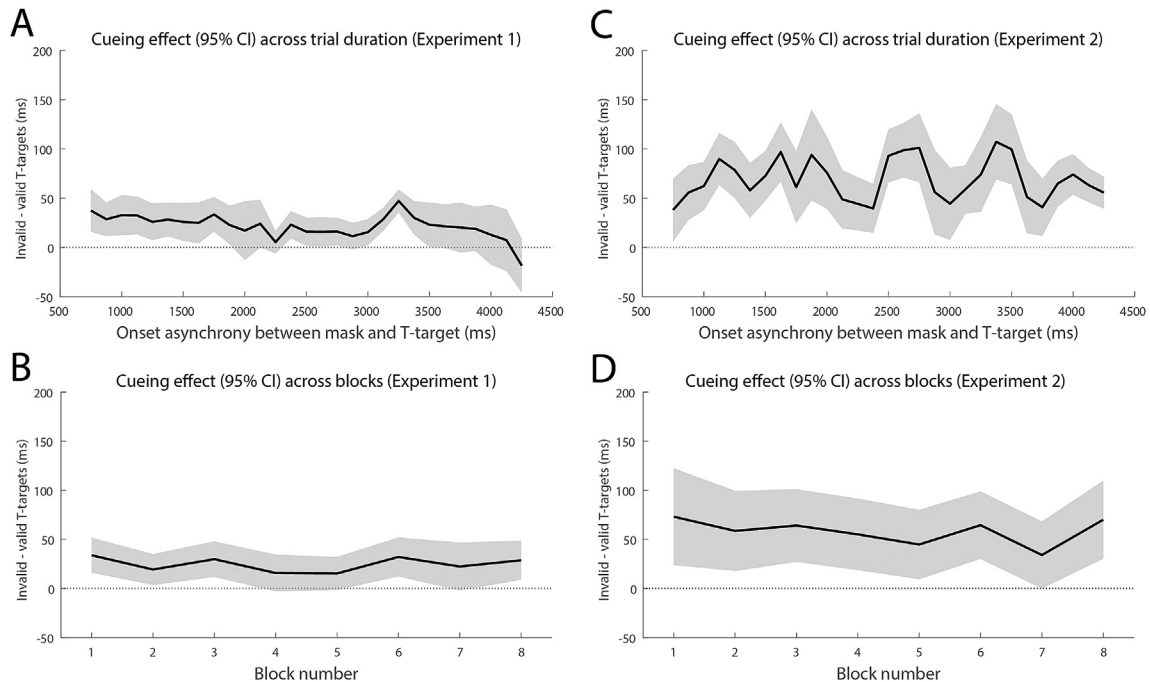


Fig. 5 – Cueing effect (response time to invalidly cued minus validly cued T-targets) over time, providing insights into the robustness of our manipulation of endogenous attention. Panels A and C depict the magnitude of the cueing effect over the course of the 5 sec trial duration for Experiments 1 and 2, respectively. T-targets could be presented between 500 msec and 4500 msec after mask onset. The black line is the average cueing effect within a 500 msec sliding window, and the gray area depicts the 95% confidence interval. Panels B and D depict the magnitude of the cueing effect across each of the 8 experimental blocks for Experiments 1 and 2, respectively. Each block comprised 24 validly cued and 8 invalidly cued T-targets. The black line is the average cueing effect within each block, and the gray area depicts the 95% confidence interval.

effect over the course of the experiment. The cueing effect was numerically larger in the first half of the experiment ($M = 73$ msec, $SD = 84$) than in the second half ($M = 67$ msec, $SD = 60$), but this difference was not reliable, as a Bayesian t -test provides 5 times more evidence for a null effect than for a difference in cueing effect between experimental halves, $BF_{10} = .20$.

Taken together, these data show that the predictive cues were successful in manipulating participants' attention, and consistently enabled participants to attend the cued color, throughout the entire trial and over the course of the entire experiment.

5.3. The influence of attention on suppression durations

Now that the manipulation of attention through predictive central cueing is established to be successful and robust across time, we can distinguish between suppressed probes that were either of the attended (i.e., cued) color or of the unattended (uncued) color (Fig. 4B). Contrary to our expectations, but similarly to Experiment 1, detection times to suppressed probes did not reliably differ depending on whether they were of the attended color ($M = 2172$ msec, $SD = 958$), or the unattended color ($M = 2191$ msec, $SD = 954$). A Bayesian t -test confirmed that the observed data was about 5 times as likely to support the null hypothesis (no difference between conditions) than it was to support the alternative hypothesis

(a difference between conditions), $BF_{10} = .21$. The evolution of the BF after addition of each participant (Fig. 4B, right panel) again shows that this null effect is robust to the addition of further participants and is even more pronounced when a less informative prior is chosen.

In order to assess whether the effect of feature-based attention on probe detection depended on individual participants' tendency to utilize the central cue, we performed a Bayesian pairwise correlation analyses between the response time difference on T-targets (valid vs invalid, reflecting the success of the attention manipulation) and the response time difference to the suppressed probes (attended vs unattended; Fig. 4C). This analysis revealed that the success of the attention manipulation was not related to the influence of feature-based attention on the detection of suppressed probes; the observed data were about 3 times more likely to support the null hypothesis (no correlation) than to support the existence of this correlation, $R = .15$, $BF_{10} = .31$.

Finally, the hit-rates for attended ($M = 76.7\%$, $SD = 17.4$) and unattended probes ($M = 76.2\%$, $SD = 18.0$) did not differ ($BF_{10} = .2$, or 4.5 times more evidence for the null than for the alternative hypothesis), showing that sustained attention did not impact the hit-rate to suppressed probes. In addition, the very low rate of false alarms (.8%, $SD = .9$) showed that observers did not randomly report probes when there were none, but actually reported the appearance of a probe when it was there. Taken together, we conclude that the observed null

effect in the response time data was (1) not undermined by an advantage for attended probes hidden in the response choice data, (2) not caused by ceiling performance in detecting the probes, as participants failed to detect about 20% of the probes on average, and (3) not caused by an overly liberal response tendency to report probe visibility. Taken together, these data demonstrate that attending to a particular color is not sufficient for enhancing conscious access of (interocularly) suppressed visual input of that color.

5.4. Comparing the influence of feature-based and spatial attention

Finally, we compared the data obtained in Experiments 1 and 2, using Bayesian independent-samples t-tests. The manipulation of feature-based attention (Experiment 2) elicited a stronger validity effect on the T-target than the manipulation of spatial attention (Experiment 1), $BF_{10} = 20$. In addition, participants' accuracy on the T-target discrimination task was more influenced by cue validity in Experiment 2 (manipulating feature-based attention) than in Experiment 1 (manipulating spatial attention), $BF_{10} = 2743$. From this, it seems that (at least within the current experimental paradigm) endogenous cues are more potent in modulating feature-based attention than spatial attention. Nonetheless, the influence of the cue on the detection of interocularly suppressed probes did not differ between Experiments, as the data were almost 4 times as likely to support a null effect than it was to support a difference between experiments, $BF_{10} = .28$.

6. Experiment 3: methods

6.1. Rationale

Experiment 2 demonstrated that even when feature-based attention is demonstrably sustained over a prolonged period, this does not necessarily cause accelerated conscious access of concurrent visual input comprising the attended feature. Based on this finding we hypothesized that observers might have used a more specific search template (a red/blue inverted T-shape), rather than two generic ones (red/blue, and inverted T-shape), such that the attentional template created for the T-target discrimination task did not impact visual processing of the initially suppressed circular probe, even when it happened to be of the cued color. If this hypothesis was true, a more generic template (say, the color red) *should* have biased conscious access of template-matching visual input.

Experiment 3 was designed to test exactly that: does endogenous attention bias conscious access when observers use a more generic attentional template, which encompasses the suppressed visual input? To address this question we used the exact same stimulation as in Experiment 2, but changed the instructions: instead of reporting the location of the inverted T-target (which was likely to be of the cued color), we now required participants to report the location of the cued (i.e., red or blue) T-target, irrespective of its specific shape (upright or inverted T-target). As such, participants were now looking for an ill-defined shape of a particular (i.e., cued) color,

thus broadening the attentional template. As in Experiments 1 and 2, participants were also required to report the appearance of a colored disc (the probe).

If participants in Experiment 1 (and 2) indeed used a specific search template (i.e., a red/blue inverted T-shape), which failed to encompass the visual characteristics of the suppressed (circular) probe, we expected that the central cues would affect conscious access to the probes now that participants were instructed to search for a red/blue target, rather than for an inverted T-shaped target that was likely to be red/blue. In contrast, if our paradigm did not allow for feature-based attention to transfer from task-relevant T-shaped targets to task-irrelevant suppressed probes, we expected not to find an influence of the central cues on the suppressed probes in Experiment 3 either.

6.2. Participants

We gathered 20 participants following the same procedure as in Experiments 1 and 2, but now from the Radboud University subject pool. Participant collection was stopped after the planned minimum of 20 participants, because the optional stopping condition ($BF > 6$) had been met. The eventual group of participants comprised 7 males with a mean age of 22.7 years ($SD = 3.1$). All participants had (corrected to) normal vision, including stereopsis (tested with the TNO test for stereoscopic vision, 12th edition; [Walraven, 1972](#)) and color vision (tested with the Ishihara color blindness test plates; [Ishihara, 1960](#)).

6.3. Additional methods, stimuli, and procedure

The first trial type (256 trials) was designed to manipulate participants' feature-based attention towards either the color red or the color blue ([Fig. 3](#), right panel). Visual stimulation was identical to that of Experiment 2 ([Fig. 3](#)). In this case, however, there were no invalid trials, as observers were instructed to report the location (left or right hemi-field) of the cued (i.e., red or blue) T-target as fast as possible, irrespective of its shape (i.e., upright or inverted T-shape).

The second trial type (64 trials) was designed to measure the influence of feature-based attention on suppression durations ([Fig. 3](#), left panel). As in Experiment 2, participants were required to provide a speeded response upon probe detection (spacebar), and probes of the attended and unattended color were presented with equal prevalence (50%, or 32 trials each).

Stimulus presentation and data collection were performed on a PC with an Intel Xeon CPU E5-1620 3.7 GHz processor, 8 Gb of RAM memory, running Windows 7 (64-bit). Responses were registered via a Windows keyboard (Corsair), and stimuli were presented on a 24" BenQ XL 2420Z LED monitor with a 144 Hz refresh rate, at the native resolution of 1920×1080 pixels. A chinrest ensured a fixed viewing distance to the screen of 95 cm. Luminance output of all stimuli was equated as much as possible with that of Experiment 2 (but was on average about 15% brighter, due to differences in the brightness scaling of the monitors). The retinal size of all stimuli, the timing of stimulus presentation, and the experimental procedure were identical to that of Experiment 2.

7. Experiment 3: results

7.1. Participants and data exclusion

One participant was excluded from further analyses because she failed to report the presence of a suppressed probe within the 5 sec trial duration on more than 50% of the trials (hit-rate of 14%). The remaining 19 participants had an average hit-rate of 98.1% (SD = 2.7) for detecting the suppressed probe.

Occasionally, observers erroneously pressed the arrow keys (used for the T-target discrimination task) when the suppressed probe was presented. RTs to the suppressed probes in those trials were excluded from further analysis, as these responses reflect a violation of the task instructions. As a result of this, the computation of the median response time to probes in the attended condition was based on 25.9 trials (SD = 3.9), of which 25.4 (SD = 3.6) were hits and therefore yielded an actual RT value. In the unattended condition, the computation of the median response times to probes was based on 30.4 trials (SD = 1.8), of which 29.7 (SD = 2.1) were hits and therefore yielded an actual RT value. No responses were provided within 250 msec after probe onset, and therefore no RTs were excluded on the basis of this exclusion rule.

As in Experiments 1 and 2, Kolmogorov–Smirnov tests indicated that the response time data violated the assumption of normality in all conditions (all p 's < .05). In contrast to Experiments 1 and 2, however, the assumption of normality was still violated after \log_e -transformation of raw response times, $K = .48$, $p < .05$. Because the RT distribution looked roughly normal (symmetry around the peak value), and to promote comparison with the findings of Experiments 1 and 2, we conducted Bayesian t -tests despite violation of the assumption of normality (but non-parametric frequentist analyses yield the same pattern of findings). Again, all response times depicted in graphs and reported in text reflect the raw response times (in seconds and milliseconds) to facilitate interpretation.

Participants were 92.3% (SD = 5.9) accurate in locating the T-target of the cued color, with an average response time of 440 msec (SD = 66).

7.2. The influence of attention on suppression durations

In contrast with the results of Experiments 1 and 2, detection times to suppressed probes were reliably faster when they were of the attended color ($M = 1132$ msec, $SD = 288$) compared to the unattended color ($M = 1204$ msec, $SD = 277$), and this was numerically observed for all but 3 participants (Fig. 6, left panel). A Bayesian t -test confirmed that the observed data were about 24 times more likely to support the alternative hypothesis (a difference between conditions) than it was to support the null hypothesis (no difference between conditions), $BF_{10} = 23.7$. The evolution of the BF after addition of each participant (Fig. 6, right panel) shows that this effect was reliable from the 9th participants onwards, and is robust to the choice of prior width.

Finally, the hit-rates for attended ($M = 98.7\%$, $SD = 2.6$) and unattended probes ($M = 97.5\%$, $SD = 3.5$) did not reliably differ ($BF_{10} = .8$). This could reflect a ceiling effect, considering the

overall high hit-rate and fast response times (perhaps caused by the higher luminance emitted by the stimuli in Experiment 3 compared to Experiments 1 and 2). As in Experiments 1 and 2, the very low rate of false alarms (.7%, $SD = .7$) suggested that observers did not randomly report probes when there were none, but actually reported the appearance of a probe when it was there.

Taken together, we conclude that participants were faster at reporting the presence of initially suppressed probes when they were of the attended (i.e., cued) color. Although we cannot assert whether this reflects faster detection of, or a more liberal decision threshold for, probes of an attended color, these data at least show that central cues (which are only informative for the T-target discrimination task) can affect response times to unrelated probe stimuli within our experimental set-up. As such the null effect observed in Experiment 2 did not reflect a general insensitivity of our experimental paradigm, but genuinely demonstrated that attending to a particular color is not a sufficient condition for accelerating conscious access of suppressed visual input of that color.

8. General discussion

Humans are dynamic agents that heavily rely on visual information to navigate through their environment. Considering that only part of the visual input gives rise to a conscious experience, a mechanism is needed to favor relevant over irrelevant visual input. Here we investigated whether endogenous attention accelerates conscious access of non-conscious visual input. The current dual-task paradigm allowed us to ensure that participants sustained spatial (Experiment 1) and feature-based (Experiment 2) attention throughout the entire period of a trial. Although the deployment of both feature-based and spatial attention reliably improved performance on the speeded discrimination of consciously accessible visual input (throughout the entire trial, and over the course of the entire experiment), neither impacted the time at which initially suppressed visual input reached conscious access. The Bayesian approach allowed us to establish that this finding reflects a genuine null effect, rather than experimental insensitivity. This shows that directing endogenous attention (to a location or color) is not sufficient for accelerating conscious access of visual input at that location, or constituted of that color.

Because this finding was relatively unexpected, we questioned whether some aspect of our design intrinsically prohibited the allocation of attention devoted to perform the T-target discrimination task, to transfer to the suppressed probes. If this was the case, top-down attention never got a fair chance to modulate conscious access in our Experiments 1 and 2. For instance, it could be an issue that (A) the cues were not informative of the suppressed probes per se, (B) that the suppressed probes were presented only rarely, (C) that the suppressed probes were ramped up in intensity rather than presented with abrupt onset, or (D) that the visual characteristics of the probes were too different from those of the targets. As such, Experiment 3 was set out to test whether our stimuli in principle allowed for the influence of central cues

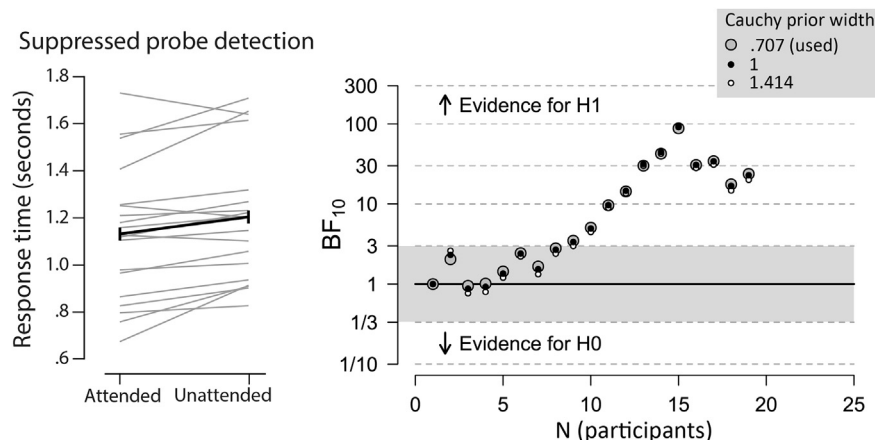


Fig. 6 – Results of Experiment 3: accelerated conscious access for probes of the attended color. This figure relates to the detection times of initially interocularly suppressed probes (in seconds), which were either of the attended (i.e., cued) or the unattended (i.e., uncued) color. The left part shows response times in seconds, with individual participants in gray, and the group average in black. The right part shows the evolution of the BF for the difference in response times between trials in which the suppressed probe was of the unattended color and trials in which the suppressed probe was of the attended color. Error bars in the response time graph depict 95% confidence intervals after removal of the between-subject variance (Cousineau, 2005; using the correction proposed by Morey, 2008). Shaded areas in the sequential analysis graph correspond to BFs that convey insufficient evidence to support either the null or the alternative hypothesis (Dienes, 2011; Jeffreys, 1961; Lee & Wagenmakers, 2014). Statistical analyses were conducted on \log_e -transformed response times (Bush et al., 1993; Gayet & Stein, 2017). Sequential analysis plots were adapted from JASP 0.8.2.0 output (JASP Team, 2017).

on the T-target discrimination task to affect detection times on the suppressed probes. Experiment 3 used the exact same visual stimulation as Experiment 2, but the task instructions were modified such as to maximize the chance that the cues would affect response times to the probes: here, the cue instructed participants to search for the red or blue red target, irrespective of its exact shape (upright or inverted), while a (red or blue) probe could be presented occasionally. In this experiment, suppressed probes were detected faster when they were of the attended color compared to when they were not. Thus, Experiment 3 showed that nothing in the visual stimulation of Experiment 2 prohibited our manipulation of top-down attention to transfer to the probe detection task. Given that (A) the visual stimulation of Experiment 2 allowed for the manipulation of attention to transfer to the probe detection task, and that (B) the manipulation of attention in Experiment 2 was successful, but yet (C) did not affect probe detection, we conclude that directing top-down attention to a color is not a sufficient condition to accelerate conscious access visual input.

Directing attention to a particular location in space is synonymous to enhancing visual perception at that location. In the case of Experiment 1, this might entail that attending to the left hemi-field did, in fact, enhance visual processing of the suppressed probes, but that the potency of the masks to suppress the probes at that location was simultaneously enhanced. The net result, following this hypothesis, is that spatial attention had no influence on the time it took to detect a suppressed probe, because processing of both the mask and the probe was enhanced. We argue, however, that even in this scenario probes should be detected faster at the attended compared to the unattended location. This follows from

Levelt's fourth proposition, which states that increasing the (in this case: effective) strength of the stimuli presented to both eyes will increase the propensity of switches in dominance between the two eyes' input (Levelt, 1965). Thus, even if it were true that the cue enhanced both the mask and the probe in our experiment, directing attention to a location should reduce interocular suppression durations for probes presented at that location. In line with this, exogenous cueing to a location in space prompts a reversal in ocular dominance at the cued location in a binocular rivalry paradigm (Paffen & Van der Stigchel, 2010). Perhaps, the greater suppression depth elicited by CFS compared to binocular rivalry suppression (Tsuchiya, Koch, Gillroy, & Blake, 2006) attenuates the attentional modulation of the suppressed probes or, along similar lines, the high-contrast CFS masks are only minimally prone to attentional (up)modulation. Alternatively, the modulation of physical stimulus strength (as in Levelt, 1965) is more similar to the modulation of effective stimulus strength through exogenous attention (as in Paffen & Van der Stigchel, 2010) than to modulation of effective stimulus strength through endogenous attention (as in our study). That is, exogenous spatial attention is known to affect both contrast and response gain, whereas endogenous spatial attention only affects response gain (Ling & Carrasco, 2006). Thus, it is not clear whether Levelt's fourth proposition applies to changes in effective stimulus strength induced by endogenous spatial attention. Irrespective of the underlying reason, spatial cueing had no influence on the detection of initially interocularly suppressed probes in our study. Hence, the deployment of sustained spatial attention was not a sufficient condition for accelerating conscious access of suppressed visual input at that location. This shows that changes in

location-specific attentional states are not necessarily accompanied by changes in conscious states.

In Experiment 2, observers attended a color. Considering that the mask was achromatic, attending (say) the color red should selectively impact processing of red probes, but not of the masks presented at the same retinal location as the probes (i.e., contra-ocularly). Nonetheless, and similar to spatial attention, attending a particular color did not accelerate conscious access of initially suppressed probes of that color. This finding is especially remarkable, considering the general consensus that feature-based attention relies upon visual working memory to keep the template available during visual search (Bundesen et al., 2005; de Fockert et al., 2001; Gunseli, Olivers, & Meeter, 2014; Wolfe, 1994), and that maintaining a color in visual working memory consistently reduces suppression times of matching probes in a b-CFS paradigm (Gayet, Paffen, et al., 2016, 2013; van Moorselaar et al., 2018). How is it possible then, that maintaining (say) the color red in visual working memory accelerates conscious access of concurrent visual input that happens to be red (e.g., Gayet, Paffen, et al., 2016, 2013; van Moorselaar et al., 2018), whereas an attentional template of the color red did not? This is especially remarkable considering that Experiment 2 comprised more participants and more trials per condition than the aforementioned studies using memory templates (i.e., between 1.5 and five times more trials in total), and should thus have a higher power to detect influences of attentional templates on conscious access. We consider three possibilities that could explain the different findings obtained with the present approach, and with direct manipulations of visual working memory content.

One possibility is that, in the current study, observers were not maintaining their attentional template in visual working memory. One notable difference between the current approach and studies that aim at manipulating the content of visual working memory is that these latter studies include many different template variations (e.g., multiple spatial locations, or color variations). As a consequence, observers in these studies need to rely on effortful working memory maintenance, as opposed to less resource-costly storage mechanisms that can be applied to overlearned stimuli (Shiffrin & Schneider, 1977; Woodman, Carlisle, & Reinhart, 2013). Indeed, when the same memory item is repeatedly memorized, its influence on concurrent attentional capture (van Moorselaar, Theeuwes, & Olivers, 2016), visual search (Gunseli, Olivers, & Meeter, 2016), and b-CFS (Gayet, van Moorselaar, Paffen, Olivers, & Van der Stigchel, submitted for publication) is abolished after a handful of repetitions. Also, when the different memoranda are dissociable at a categorical rather than within-category level (e.g., red vs blue, rather than two variations of blue) attentional capture by memory-matching stimuli is abolished as well (Experiment 2 of Olivers, Meijer, & Theeuwes, 2006). In both cases (categorical and repeated memoranda) one could hypothesize that observers no longer have to rely on visual working memory but could rely on (verbal-labels in) more long-term (and less visual) memory, which has a lesser impact on concurrent processing of visual input. In the current study, however, despite using only two colors and two locations, the attentional template did reliably affect the response speed to

consciously accessible (i.e., to-be-searched-for) targets. As such, these issues cannot be responsible for the absence of an influence of sustained attention on suppression times in.

This leads to the second possibility: attention does not operate on non-conscious (or more specifically: interocularly suppressed) information. Studies using binocular rivalry have shown that sustained endogenous attention can modulate interocular suppression (Chong, Tadin, & Blake, 2005; Mitchell, Stoner, & Reynolds, 2004; Ooi & He, 1999; Zhang, Jiang, & He, 2012; for reviews on attentional modulation of binocular rivalry see, Paffen & Alais, 2011; Dieter & Tadin, 2011). In those studies, however, endogenous attention affected suppression durations by prolonging perception of the dominant (i.e., consciously accessible) image, rather than shortening perception of the non-dominant (i.e., non-conscious) image. As such, it remains unclear whether or not endogenous attention impacts conscious access of initially non-conscious visual input. Using CFS improves the experimenter's control over which of two images is suppressed, so that the influence of sustained attention on non-conscious visual input can be isolated from that on conscious visual input. Using CFS, Kanai et al. (2006) showed that feature-based attention enhanced visual processing of suppressed gratings, causing them to elicit a more pronounced after-effect (also, see Melcher et al., 2005). This is in contrast with our current findings. One could argue that feature-based attention suffices for enhancing the processing of non-conscious visual input, but does not suffice for rendering non-conscious visual input consciously accessible. This interpretation seems unlikely, however, as an enhancement of non-consciously accessible visual input should by definition lower its effective threshold for reaching conscious access (Gayet, Paffen, et al., 2016).

A third possibility then pertains to a critical difference between our study and that of Kanai et al. (2006): The CFS suppressed stimuli in the study of Kanai and colleagues were identical to the visible (i.e., to-be-attended) targets, whereas this was not the case in the present study. Rather, participants in our studies were deploying attentional resources to detect a specifically shaped stimulus (i.e., an inverted T-shape), while the suppressed probe was of a different shape (circular). As such, it is possible that participants' attentional bias towards (e.g., red) inverted T-shaped targets did not generalize to differently shaped targets, such as our suppressed circle-shaped probe. This explanation is particularly compelling, considering that we did observe an effect of feature-based attention on the suppressed probes when participants were cued to search for a red/blue target of an ill-defined shape (i.e., an upright or inverted T-shape), inciting them to instate a broader attentional template (Experiment 3). From this, we tentatively conclude that observers can deploy attentional resources to act upon a specific contingency of features (e.g., location or color and shape), such that conscious access is only favored for visual input that matches this specific contingency of features.

How does this putative specificity of attentional templates relate to what is known about memory templates? Earlier studies have shown that maintaining a color in working memory enhances conscious access to all visual input comprising that color, irrespective of its shape (e.g., Experiments 1 & 5 of Gayet et al., 2013), whereas directing top-down

attention to a color does not enhance conscious access to all visual input comprising that color (current Experiment 2). This is unexpected, considering that the sustained deployment of top-down attention (e.g., to a location or color) requires the recruitment of (spatial/visual) working memory for keeping the attended location or color available. Recent evidence suggests that the representation maintained in working memory (i.e., the template) can differ depending on the instruction to either *search* for a particular feature (and thereby to memorize it), or to explicitly *memorize* a feature (Gunselli, Meeter, et al., 2014), with the former comprising a stronger top-down component (van Driel, Gunselli, Meeter, & Olivers, 2017). As such, a difference in task instruction could determine the specific content that is loaded into working memory: when instructed to memorize the color red (*memorization* instruction), participants will memorize the color red, which will affect the processing of all subsequent red visual input. In contrast, when instructed to search for a particular red target (*search* instruction), participants might memorize the color red in conjunction with the shape of the to-be-searched-for target. According to this view, observers in our Experiment 2 (using a *search* instruction) created an attentional template representing a red inverted T-shaped target, thereby favoring subsequent red inverted T-shaped targets, but not necessarily other red (non-target) objects. Indeed, in the case of explicit *memorization* instructions, only the to-be-remembered feature dimension of a memorized object elicits a sustained content-specific neural trace in visual processing areas (Serences, Ester, Vogel, & Awh, 2009). Consequently, only the to-be-remembered feature dimension of a memorized object enhances conscious access of matching visual input (Experiment 5 of Gayet et al., 2013; Experiment 3 of Gayet, Brascamp, Van der Stigchel, & Paffen, 2015), and attentional capture by matching visual input (e.g., Experiment 4 of Olivers et al., 2006). In contrast, when observers are instructed to *search* for a particular target shape (independent of its color), as cued by a colored shape, search times are reduced when the eventual (shape-matching) target also matches the incidental color of the cued shape (Foerster & Schneider, 2018). This reveals that, unlike the case of a *memorization* instruction, in case of a *search* instruction, the attentional template can encompass more feature dimensions than only the feature dimension that is explicitly cued. Considering that maintaining two features of an object in visual working memory does not drain substantially more resources than maintaining a single feature in visual working memory (Luck & Vogel, 1997), it is resource-efficient to utilize an attentional template comprised of a conjunction of two features if this potentially benefits task performance.

Let us consider the hypothesized situation that participants in our Experiment 2 (*search* instruction) did indeed load a conjunction of a red T-shaped target into visual working memory; their mnemonic template should still have enhanced any visual input that is red, albeit to a lesser extent than visual input that is red as well as T-shaped (e.g., Foerster & Schneider, 2018). In fact, attending a feature can even enhance consciously accessible visual input that shares very little perceptual similarity with the attentional template, aside from the attended feature (and even when presented at an irrelevant spatial location; White & Carrasco, 2011), and our Experiment 3

showed that attending a feature can also accelerate detection of initially non-conscious visual input (using the exact same visual stimulation as in Experiment 2). So why would an attentional template composed of a conjunction of multiple features (e.g., a red inverted T-shape) not accelerate conscious access of visual input that comprises any of those features (e.g., a red circle)? The discrepancy between our current findings and the traditional view that attention indiscriminately biases visual selection in favor of the attended feature is reminiscent of the view proposed by Rafal et al. (2002). It has indeed been shown often that attention to simple features (e.g., color, shape, etc.) can bias visual selection by instating a bias towards the attended feature in early cortical areas representing that feature (e.g., Desimone & Duncan, 1995). In case the observer's action depends on a more complex conjunction of features (higher-level analysis of words in the case of Rafal et al., or color-shape conjunctions in the present case), however, attention can gate visual selection at a higher level of representation, which is used for action selection. Following Rafal et al.'s view, this suggests that while observers in Experiment 2 were attending red/blue (as part of a red/blue inverted T-shaped target), the *shape* of the attentional template was relevant for the response (as observers were instructed to report the location of the inverted T-shape), whereas its color was not (as the target could also be of the other color). The findings of Experiment 3 offer support to this view. Here, the instructions were tweaked such that the action-relevant part of the attentional template was the *color* rather than the *shape*; in this case we did observe faster detection of suppressed probes when they were of the cued (i.e., attended) color.

9. Conclusion

Further research is needed to map out exactly what boundary conditions have to be met for attention to accelerate conscious access of suppressed visual input. At present, we conclude that deploying attentional resources to facilitate detection of stimuli of a particular color (at the expense of differently colored stimuli) is not a sufficient condition for accelerating conscious access of suppressed visual input that is incidentally constituted of the same color. Similarly, deploying attentional resources to facilitate detection of a stimulus at a specific location is not a sufficient condition for accelerating conscious access of suppressed visual input that is incidentally positioned at that same location. In light of the relation between consciousness and attention, the current data – at least – oppose theories of consciousness that equate consciousness with attention: attention to a color or location was measurably modulated, but this did not necessarily affect the content of consciousness. It is unclear from our data, however, whether attention is necessary for consciousness to emerge, or whether attentional orienting is dependent on consciousness. That is, endogenous attention might have impacted the processing of non-conscious visual input, without affecting conscious access to this visual input. At the same time, we showed that endogenous attention can affect conscious access, under certain specific constraints. Possibly, when attentional templates comprise multiple features (e.g., a color and a shape) only that feature which is required for the

observer to respond will affect conscious access. This is in line with the view of Rafal et al. (2002), who showed that attention selectively gates conscious access when it is of benefit to the behavior of the observer in neglect patients. This was not the case in our study (Experiments 1 and 2), as the color or location of the suppressed probes were irrelevant to the observers' task at hand (i.e., pertaining to the location or orientation of the target). What the present findings do show is that directing endogenous attention is not a sufficient condition for consciousness to emerge.

Funding

This project has received funding from the Netherlands Organization for Scientific Research (NWO VIDI Grant 45213008 to Stefan Van der Stigchel).

Declarations of interests

None.

Acknowledgements

The authors thank Micha Heilbron for programming and piloting an earlier version of the experiments presented here, and Marius Peelen for providing feedback on the revised manuscript and the design of Experiment 3.

REFERENCES

- Awh, E., & Jonides, J. (2001). Overlapping mechanisms of attention and spatial working memory. *Trends in Cognitive Sciences*, 5(3), 119–126.
- Baars, B. J. (2005). Global workspace theory of consciousness: Toward a cognitive neuroscience of human experience. *Progress in Brain Research*, 150, 45–53.
- Block, N. (2011). Perceptual consciousness overflows cognitive access. *Trends in Cognitive Sciences*, 15(12), 567–575.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10, 433–436.
- Bundesden, C., Habekost, T., & Kyllingsbæk, S. (2005). A neural theory of visual attention: Bridging cognition and neurophysiology. *Psychological Review*, 112(2), 291.
- Buracas, G. T., & Boynton, G. M. (2007). The effect of spatial attention on contrast response functions in human visual cortex. *Journal of Neuroscience*, 27(1), 93–97.
- Bush, L. K., Hess, U., & Wolford, G. (1993). Transformations for within-subject designs: A Monte Carlo investigation. *Psychological Bulletin*, 113(3), 566–579.
- Chong, S. C., Tadin, D., & Blake, R. (2005). Endogenous attention prolongs dominance durations in binocular rivalry. *Journal of Vision*, 5(11), 6.
- Cousineau, D. (2005). Confidence intervals in within-subject designs: A simpler solution to Loftus and Masson's method. *Tutorials in Quantitative Methods for Psychology*, 1(1), 42–45.
- Dehaene, S., Changeux, J. P., Naccache, L., Sackur, J., & Sergent, C. (2006). Conscious, preconscious, and subliminal processing: A testable taxonomy. *Trends in Cognitive Sciences*, 10(5), 204–211.
- Dehaene, S., Kerszberg, M., & Changeux, J. P. (1998). A neuronal model of a global workspace in effortful cognitive tasks. *Proceedings of the National Academy of Sciences*, 95(24), 14529–14534.
- Dennett, D. C. (1991). *Consciousness explained*. Boston, MA: Little, Brown and Co.
- Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, 18(1), 193–222.
- Dienes, Z. (2011). Bayesian versus orthodox statistics: Which side are you on? *Perspectives on Psychological Science*, 6(3), 274–290.
- Dieter, K. C., & Tadin, D. (2011). Understanding attentional modulation of binocular rivalry: A framework based on biased competition. *Frontiers in Human Neuroscience*, 5, 155.
- van Driel, J., Gunseli, E., Meeter, M., & Olivers, C. N. (2017). Local and interregional alpha EEG dynamics dissociate between memory for search and memory for recognition. *NeuroImage*, 149, 114–128.
- Edelman, G., & Tononi, G. (2008). *A universe of consciousness: How matter becomes imagination*. New York, NY: Basic books.
- de Fockert, J. W., Rees, G., Frith, C. D., & Lavie, N. (2001). The role of working memory in visual selective attention. *Science*, 291(5509), 1803–1806.
- Foerster, R., & Schneider, W. X. (2018). Involuntary top-down control by search-irrelevant features: Visual working memory biases attention in an object-based manner. *Cognition*, 172, 37–45.
- Gandhi, S. P., Heeger, D. J., & Boynton, G. M. (1999). Spatial attention affects brain activity in human primary visual cortex. *Proceedings of the National Academy of Sciences*, 96(6), 3314–3319.
- Gayet, S., Brascamp, J. W., Van der Stigchel, S., & Paffen, C. L. E. (2015). Cogito ergo video: Task-relevant information is involuntarily boosted into awareness. *Journal of Vision*, 15(5), 3.
- Gayet, S., Guggenmos, M., Christophel, T. B., Haynes, J.-D., Paffen, C. L. E., Van der Stigchel, S., et al. (2017). Visual working memory enhances the neural response to matching visual input. *Journal of Neuroscience*, 37(28), 6638–6647.
- Gayet, S., Paffen, C. L., Belopolsky, A. V., Theeuwes, J., & Van der Stigchel, S. (2016). Visual input signaling threat gains preferential access to awareness in a breaking continuous flash suppression paradigm. *Cognition*, 149, 77–83.
- Gayet, S., Paffen, C. L. E., & Van der Stigchel, S. (2013). Information matching the content of visual working memory is prioritized for conscious access. *Psychological Science*, 24(12), 2472–2480.
- Gayet, S., & Stein, T. (2017). Between-subject variability in the breaking continuous flash suppression paradigm: Potential causes, consequences, and solutions. *Frontiers in Psychology: Consciousness Research*, 8, 437.
- Gayet, S., van Maanen, L., Heilbron, M., Paffen, C. L. E., & Van der Stigchel, S. (2016). Visual input that matches the content of visual working memory requires less (not faster) evidence sampling to reach conscious access. *Journal of Vision*, 16(11), 26.
- Gayet, S., Van Moorselaar D., Paffen C.L.E., Olivers C.N.L., & Van der Stigchel, S. (2018). Transiting information from short- to long-term memory reduces the perceptual bias towards matching visual input, under revision.
- Gayet, S., Van der Stigchel, S., & Paffen, C. L. E. (2014a). Seeing is believing: Utilization of subliminal symbols requires a visible relevant context. *Attention, Perception, & Psychophysics*, 76, 489–507.
- Gayet, S., Van der Stigchel, S., & Paffen, C. L. E. (2014b). Breaking continuous flash suppression: Competing for consciousness on the pre-semantic battlefield. *Frontiers in Psychology: Consciousness Research*, 5, 460.
- Gilroy, L. A., & Blake, R. (2005). The interaction between binocular rivalry and negative afterimages. *Current Biology*, 15(19), 1740–1744.
- Graziano, M. S. (2013). *Consciousness and the social brain*. Oxford, UK: Oxford University Press.

- Greenberg, A. S., Esterman, M., Wilson, D., Serences, J. T., & Yantis, S. (2010). Control of spatial and feature-based attention in frontoparietal cortex. *Journal of Neuroscience*, 30(43), 14330–14339.
- Gunseli, E., Meeter, M., & Olivers, C. N. (2014). Is a search template an ordinary working memory? Comparing electrophysiological markers of working memory maintenance for visual search and recognition. *Neuropsychologia*, 60, 29–38.
- Gunseli, E., Olivers, C. N., & Meeter, M. (2014). Effects of search difficulty on the selection, maintenance, and learning of attentional templates. *Journal of Cognitive Neuroscience*, 26(9), 2042–2054.
- Gunseli, E., Olivers, C. N., & Meeter, M. (2016). Task-irrelevant memories rapidly gain attentional control with learning. *Journal of Experimental Psychology: Human Perception and Performance*, 42(3), 354–362.
- Hassin, R. R. (2013). Yes it can: On the functional abilities of the human unconscious. *Perspectives on Psychological Science*, 8(2), 195–207.
- Hesselmann, G., & Moors, P. (2015). Definitely maybe: Can unconscious processes perform the same functions as conscious processes? *Frontiers in Psychology*, 6, 584.
- Hofstoetter, C., Koch, C., & Kiper, D. C. (2004). Motion-induced blindness does not affect the formation of negative afterimages. *Consciousness and Cognition*, 13(4), 691–708.
- Hsu, S. M., George, N., Wyart, V., & Tallon-Baudry, C. (2011). Voluntary and involuntary spatial attentions interact differently with awareness. *Neuropsychologia*, 49(9), 2465–2474.
- Ishihara, S. (1960). *Tests for colour-blindness*. Tokyo, Japan: Kanehara Shuppan Company.
- James, W. (1890/1950). *The principles of psychology*. New York, NY: Dover Publications.
- JASP Team. (2017). *JASP (version 0.8.2.0) [Computer software]*.
- Jeffreys, H. (1961). *Theory of probability* (3rd ed.). Oxford, UK: Oxford University Press, Clarendon Press.
- Jiang, Y., Costello, P., & He, S. (2007). Processing of invisible stimuli: Advantage of upright faces and recognizable words in overcoming interocular suppression. *Psychological Science*, 18(4), 349–355.
- Kanai, R., Tsuchiya, N., & Verstraten, F. A. J. (2006). The scope and limits of top-down attention in unconscious visual processing. *Current Biology*, 16(23), 2332–2336.
- Kaufman, L. (1963). On the spread of suppression and binocular rivalry. *Vision Research*, 3(9), 401–415.
- Kiefer, M., & Brendel, D. (2006). Attentional modulation of unconscious “automatic” processes: Evidence from event-related potentials in a masked priming paradigm. *Journal of Cognitive Neuroscience*, 18(2), 184–198.
- Koch, C. (2004). *The quest for consciousness: A neuroscientific approach*. Atlanta, GA: Roberts & Co.
- Koch, C., & Tsuchiya, N. (2007). Attention and consciousness: Two distinct brain processes. *Trends in Cognitive Sciences*, 11(1), 16–22.
- Koch, C., & Tsuchiya, N. (2012). Attention and consciousness: Related yet different. *Trends in Cognitive Sciences*, 16(2), 103–105.
- LaBar, K. S., Gitelman, D. R., Parrish, T. B., & Mesulam, M. M. (1999). Neuroanatomic overlap of working memory and spatial attention networks: A functional MRI comparison within subjects. *NeuroImage*, 10(6), 695–704.
- Lamme, V. A. (2003). Why visual attention and awareness are different. *Trends in Cognitive Sciences*, 7(1), 12–18.
- Lee, M. D., & Wagenmakers, E. J. (2014). *Bayesian cognitive modeling: A practical course*. Cambridge, UK: Cambridge University Press.
- Levelt, W. J. (1965). *On binocular rivalry* (Doctoral dissertation). Assen: Van Gorcum.
- Ling, S., & Carrasco, M. (2006). Sustained and transient covert attention enhance the signal via different contrast response functions. *Vision Research*, 46(8), 1210–1220.
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390(6657), 279–281.
- Mack, A., & Rock, I. (1998). *Inattention blindness* (Vol. 33). Cambridge, MA: MIT press.
- Martens, U., Ansorge, U., & Kiefer, M. (2011). Controlling the unconscious: Attentional task sets modulate subliminal semantic and visuomotor processes differentially. *Psychological Science*, 22(2), 282–291.
- Maruya, K., & Blake, R. (2009). Spatial spread of interocular suppression is guided by stimulus configuration. *Perception*, 38(2), 215–231.
- Melcher, D., Pappathomas, T. V., & Vidnyánszky, Z. (2005). Implicit attentional selection of bound visual features. *Neuron*, 46(5), 723–729.
- Merrikhi, Y., Clark, K., Albarran, E., Parsa, M., Zirnsak, M., Moore, T., et al. (2017). Spatial working memory alters the efficacy of input to visual cortex. *Nature Communications*, 8, 15041.
- Mitchell, J. F., Stoner, G. R., & Reynolds, J. H. (2004). Object-based attention determines dominance in binocular rivalry. *Nature*, 429(6990), 410–413.
- Morey, R. D. (2008). Confidence intervals from normalized data: A correction to Cousineau (2005). *Reason*, 4(2), 61–64.
- van Moorselaar, D., Gayet, S., Paffen, C. L. E., Theeuwes, J., Van der Stigchel, S., & Olivers, C. N. L. (2018). Competitive interactions in visual working memory drive access to awareness. *Cortex*, 102, 6–13.
- van Moorselaar, D., Theeuwes, J., & Olivers, C. N. (2016). Learning changes the attentional status of prospective memories. *Psychonomic Bulletin & Review*, 23(5), 1483–1490.
- Naccache, L., Blandin, E., & Dehaene, S. (2002). Unconscious masked priming depends on temporal attention. *Psychological Science*, 13(5), 416–424.
- Olivers, C. N., Meijer, F., & Theeuwes, J. (2006). Feature-based memory-driven attentional capture: Visual working memory content affects visual attention. *Journal of Experimental Psychology: Human Perception and Performance*, 32(5), 1243–1265.
- Ooi, T. L., & He, Z. J. (1999). Binocular rivalry and visual awareness: The role of attention. *Perception*, 28(5), 551–574.
- Paffen, C. L. E., & Alais, D. (2011). Attentional modulation of binocular rivalry. *Frontiers in Human Neuroscience*, 5, 105.
- Paffen, C. L. E., & Van der Stigchel, S. (2010). Shifting spatial attention makes you flip: Exogenous visual attention triggers perceptual alternations during binocular rivalry. *Attention, Perception, & Psychophysics*, 72(5), 1237–1243.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10(4), 437–442.
- Posner, M. I. (1994). Attention: The mechanisms of consciousness. *Proceedings of the National Academy of Sciences*, 91(16), 7398–7403.
- Postle, B. R., Stern, E., Rosen, B. R., & Corkin, S. (2000). An fMRI investigation of cortical contributions to spatial and nonspatial visual working memory. *NeuroImage*, 11, 409–423.
- Rafal, R., Danziger, S., Grossi, G., Machado, L., & Ward, R. (2002). Visual detection is gated by attending for action: Evidence from hemispatial neglect. *Proceedings of the National Academy of Sciences*, 99(25), 16371–16375.
- Serences, J. T., Ester, E. F., Vogel, E. K., & Awh, E. (2009). Stimulus-specific delay activity in human primary visual cortex. *Psychological Science*, 20(2), 207–214.
- Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning,

- automatic attending and a general theory. *Psychological Review*, 84(2), 127–190.
- Stein, T., Hebart, M. N., & Sterzer, P. (2011). Breaking continuous flash suppression: A new measure of unconscious processing during interocular suppression? *Frontiers in Human Neuroscience*, 5, 167.
- Treue, S. (2001). Neural correlates of attention in primate visual cortex. *Trends in Neurosciences*, 24(5), 295–300.
- Treue, S., & Trujillo, J. C. M. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. *Nature*, 399(6736), 575–579.
- Tsuchiya, N., & Koch, C. (2005). Continuous flash suppression reduces negative afterimages. *Nature Neuroscience*, 8(8), 1096–1101.
- Tsuchiya, N., Koch, C., Gilroy, L. A., & Blake, R. (2006). Depth of interocular suppression associated with continuous flash suppression, flash suppression, and binocular rivalry. *Journal of Vision*, 6(10), 6.
- Walraven, J. (1972). *TNO test for stereoscopic vision*. Utrecht, The Netherlands: Lameris Instrumenten.
- Watanabe, M., Cheng, K., Murayama, Y., Ueno, K., Asamizuya, T., Tanaka, K., et al. (2011). Attention but not awareness modulates the BOLD signal in the human V1 during binocular suppression. *Science*, 334(6057), 829–831.
- White, A. L., & Carrasco, M. (2011). Feature-based attention involuntarily and simultaneously improves visual performance across locations. *Journal of Vision*, 11(6), 15.
- Wolfe, J. M. (1994). Guided search 2.0 a revised model of visual search. *Psychonomic Bulletin & Review*, 1(2), 202–238.
- Woodman, G. F., Carlisle, N. B., & Reinhart, R. M. (2013). Where do we store the memory representations that guide attention? *Journal of Vision*, 13(3), 1.
- Zhang, P., Jiang, Y., & He, S. (2012). Voluntary attention modulates processing of eye-specific visual information. *Psychological Science*, 23(3), 254–260.