Neural activity linked with visual awareness and task-relevance in a novel 2x2 design

A Thesis

Presented to

The Division of Philosophy, Religion, Psychology, and Linguistics

Reed College

In Partial Fulfillment

of the Requirements for the Degree

Bachelor of Arts

Andrew Kyroudis

December 2018

Approved for the Division

(Psychology)

Michael Pitts

Acknowledgments

Thank you friends, family, and professors for pushing me to push myself. Michael, thank you for bringing me here and keeping me here. I’m excited to see what follows.

Table of Contents

List of Tables

List of Figures

Abstract

The current study is based on previous research by Dehaene et al. (2001) in which perceptual awareness of visual stimuli was manipulated using pattern masking - a common paradigm in which salient stimuli can be made perceptually invisible by "masking stimuli" that appear in close temporal proximity. When neural activity associated with "seen" stimuli (lightly masked) was contrasted with neural activity associated with "unseen" stimuli (heavily masked), a consistent neural difference emerged: the P3b event related potential (ERP). However, the critical stimuli were always task-relevant (subjects were always reporting what they saw on each trial), and recent concerns have been raised about pre- and post-perceptual processing being mistaken for *true* conscious processing. For example, the neural activity associated with perceiving the critical stimuli could have been contaminated with neural activity associated with *reporting* about perceiving the critical stimuli.

The current two-part experiment replicates and extends Dehaene et al. (2001)'s study by including both task-relevant (trial-by-trial reporting) and task-irrelevant (no report) conditions. The current study was aimed at answering the following question: are the neural differences between seen and unseen stimuli found in task-relevant conditions also present when the same stimuli are task-irrelevant? Scalp electroencephalography (EEG) recordings and ERPs in human subjects were used to track the moment-by-moment activity of populations of neurons associated with generating conscious visual experience. The analyses focused on the P3b component, first replicating previous findings in task-relevant conditions and then testing for its presence in novel task-irrelevant conditions.

Introduction

Simply put, *consciousness* is everything we know. The scientific study of consciousness is an illuminating inquiry into the neurophysiological mechanisms from which our subjective experience emerges. As the ancient puzzlement of the mind-body problem becomes grounded in neurophysiological data, the long-standing philosophical problem of consciousness has started to become approachable by mainstream neuroscience. Starting in the early 1990s, researchers in this field advocated for simplifying the problem by first looking for consistent 'neural correlates of consciousness' (NCCs). Christof Koch, one of the first consciousness researchers in this modern wave, defines the NCC as "the minimal neuronal mechanisms jointly sufficient for any one specific conscious percept" (Koch 2004, Crick and Koch 1995). These NCCs, once identified, can then help guide future research into the detailed physiological mechanisms that constitute our conscious experience.

Because the term 'consciousness' is used in various contexts with different meanings, it is important to start by clarifying the distinction between *states* of consciousness and the *content* of consciousness. *States* of consciousness refer to states of wakefulness and non-wakefulness - for instance, the differences between normal waking life, REM sleep, and being in a coma are differences in *state* (and not necessarily *content, e.g. we can consciously experience a friend's face in waking life and consciously experience this same friend's face in our dreams*). The current investigation, on the other hand, is concerned with the *content* of consciousness. Changes in *content*, for example, are reflected in the differences between perceiving an empty green room versus an empty blue room. Within-state manipulations are used to isolate neurophysiological signals associated with changes in the *content* of consciousness, as is the standard in this line of study (Dehaene et al. 2014). 'Conscious perception' refers to first order perceptual awareness - the experience of some (visual, auditory, somatosensory) information. First order perceptual awareness is distinct from higher order perceptual awareness, where one may remember perceptual content, name the content, associate the content with other content, and so on… these phenomena are all dubbed *metacognition*. First order perceptual awareness is pure perception of content, with no extra (meta)cognition. The search for the brain processes underlying perceptual awareness is at the center of the search for the NCC (Dehaene et al. 2017).

The contrast of interest in the current study is that between neural activity associated with perceptual awareness (of certain visual stimuli) and neural activity associated with lack of perceptual awareness (of certain visual stimuli, all other things being equal). Therefore, the purpose of the current study is to help elucidate the unique properties of conscious processing from those of nonconscious processing. Simply put, the field of cognitive neuroscience in its current state has an impoverished understanding of what separates conscious from nonconscious perceptual processes. The current investigation aims to augment this incomplete neurophysiological understanding.

Why is it Useful and Interesting to Study Perceptual Awareness?

Studying the neurophysiological correlates of perceptual awareness allows us to better understand the temporal and spatial characteristics of our waking human experience. The goal is to utilize neurophysiological information to infer subjective perceptual content. For instance, when people are in a persistent vegetative state (PVS), it is impossible to tell from the "outside" if the person in PVS is alive *and* conscious or is just simply alive. Understanding the neural signatures of conscious perception can directly inform decisions about proper medical treatment for patients in unconventional states of consciousness. It is important to know if the lights are on in the patient, i.e. they are having conscious experiences, or if there is no one home and everything is dark (no experience). Currently, only positive signs of consciousness are useful (Owen et al. 2007, Owen et al. 2006) but eventually it should be possible to measure with high accuracy whether a patient is conscious or not, and either type of indicator could be used to critically inform treatment plans, improving the medical care for patients and their loved ones.

Additionally, the intersection between neural electrophysiology and technology is robust. Physically disabled patients who lack the ability to speak and to use hardware like a conventional keyboard still need to communicate with family, friends, and medical professionals. Thus, researchers have been developing methods of using electrophysiology to interact with a virtual keyboard, which allows patients to communicate linguistically, in real time, with the audience who can read the words that are being typed out via brain activity (İşcan et al. 2018, Mak et al. 2009). Some of these brain-computer interfaces utilize the same signal that is of interest in (the current) perceptual awareness research: the P3b event-related potential (ERP) (Marchetti et al. 2014). Therefore, increased insight into what this electrical brain signal correlates with could directly aid our understanding of how to use such a signal in brain-computer interfaces (BCI). Overall, there is much potential in using electrophysiology to interact with computers - specifically to mimic the functionality of biological systems that are affected by disability – and eventually this will also include BCI devices to compensate for disorders of consciousness.

The final (and arguably most interesting) reason for studying the neural correlates of perceptual awareness is that this line of research informs our understanding of the perceptual experience that *is* waking human life. How is our perceptual experience constructed from patterns of firing in neural populations? When, where, why, and how do certain brain processes produce consciousness while the vast majority of these neural processes do not? Nonconscious influences demonstrate the malleability of conscious processing. Colors, words, and more complex stimuli that are processed nonconsciously can reliably modulate conscious processing in significant ways (Naccache et al. 2001, Kouider et al. 2001, Kouider et al. 2007, Nakamura et al. 2007, Kentridge et al. 2004). Fundamentally, however, (conscious) perception is a phenomenon that relies on unconscious inferential processes that make educated guesses about perceptual content based on accumulated evidence (Jia et al. 2018, Teufel et al. 2018, Bar et al. 2004, Bar et al. 1996). In other words, conscious processing seems to be the most ‘endorsed’ outcome of nonconscious processing.

Empirical Theories of Consciousness

There are two main categories of empirical theories of consciousness: *sensory theories* posit that perception is ‘front-loaded’, meaning that perceptual processes are richly detailed, but the ability to consciously access perception is limited (by limitations in other systems, such as those supporting memory and language). *Sensory theories* assert that conscious perceptual processing takes place in sensory areas (e.g. occipito-temporal cortex for visual consciousness), while fronto-parietal areas are involved in accessing and *reporting* on conscious *content*. On the other hand, *cognitive theories* posit that sensory processing is unconscious or preconscious until this information is accessed by other, high-level systems like attention and working memory. Thus cognitive theories posit that fronto-parietal areas are necessary for conscious processing of perceptual *content*. In each main category, there are a couple of leading theories...

Cognitive Theories of Conscious Processing

The *Global Neuronal Workspace Theory (GNWT)* and *Higher Order Thought (HOT)* theory are the two main cognitive theories of consciousness. The common thread among these theories that allows them to be grouped together is the belief that perceptual awareness is reliant upon sensory and perceptual processes in conjunction with other cognitive processes.

Global Neuronal Workspace Theory

GNWT is based off of the Global Workspace Theory of consciousness, by Bernard Baars, who outlined the case for consciousness as a functional, adaptive biological phenomenon - not a common belief at the time (Dehaene et al 2011, Dehaene et al. 2006, Dehaene et al. 2004, Baars 1989). Accordingly, the modern GNWT is a functional theory of conscious processing, meaning that there is a purpose to such conscious processing - namely, widespread information availability for flexible use by cognitive systems. For GNWT, nonconscious processing becomes conscious processing precisely when information enters the global workspace. More specifically, this means that following the initial (sensory) feedforward sweep, local recurrent processing in sensory (occipito-temporal) areas occurs, followed by activity from long-range fronto-parietal neurons that result in sustained global processing - which is *exactly* when conscious processing takes place, according to this theory. As such, proponents of GNWT propose that the true neural correlates of conscious perception happen later in time, roughly in the 300-600 ms time window. This time estimate significantly differs from the temporal estimate provided by proponents of sensory theories of conscious processing, who claim that the true NCC occurs during local recurrent processing, roughly in the 100-200 ms time window. This difference in temporal prediction is a major point of contention in the study of conscious processing, and this difference is reflected in Fig 3.

According to GNWT, nonconscious processing comes in three flavors: (1) subliminal (unattended), (2) subliminal (attended), and (3) preconscious processing.

Subliminal (unattended) processing originates from weak bottom-up sensory signals that are not processed deeply and have virtually no cognitive trace; while differently, subliminal (attended) processing originates from strong bottom-up sensory signals that are (somewhat deeply) processed, potentially reaching fronto-parietal areas. Subliminal (unattended) processing results in virtually no priming capabilities - the information is quite literally "in one ear and out the other". Contrastingly, subliminal (attended) processing can induce priming effects, and is therefore strong enough to impact cognition, but is not strong enough to result in conscious experience. Both of these nonconscious processing types differ from preconscious processing, which results from strong sensory activation that is processed deeply enough to allow a multitude of priming effects, but does not lead to conscious experience due to the lack of attentional amplification. Preconscious processing is turned into conscious experience when attentional systems interact with perceptual systems and allow for global broadcasting of the pattern of perceptual information. Thus, we arrive at a crucial feature of GNWT: for conscious processing to occur, sufficient bottom-up (sensory) activation ***as well as*** sufficient top-down (attentional) amplification must occur. In other words, for conscious visual perception to occur, occipito-temporal areas must work in tandem with fronto-parietal areas in order to widely broadcast the information - resulting in visual perceptual awareness (Dehaene et al. 2006).

Higher Order Thought

Similarly, HOT supports the notion that activity in primary sensory areas (alone) is not sufficient for conscious perceptual processing. Conscious processing "crucially depends on higher-order representations," according to HOT (Lau et al. 2011). This means that conscious processing of some first order (sensory) information takes place when first order (sensory) information is simulated (or indexed) by a secondary, higher-order mechanism. Proponents of HOT believe that any cognition that includes awareness (i.e. conscious perception) is actually the experience of a metacognitive state relating to the relevant sensory content. Importantly, HOT differs from GNWT in that it does not make any predictions about the function of conscious processing (remember, GNWT assigns a function to conscious processing: sustained, global information sharing across brain areas in order to support flexible behavior). Like GNWT, though, HOT predicts that the true NCC is somewhere roughly in the 300-600 ms time window (i.e. late activity), and most likely involves neural populations in the prefrontal cortex (Lau et al. 2011).

Sensory Theories of Conscious Processing

The *Recurrent Processing (RP)* theory and *Integrated Information Theory (IIT)* are currently the two leading sensory theories of consciousness. The core belief of sensory theories is that conscious perceptual processes do not need any cognitive processing outside of perceptual processing per se. In other words, perceptual processing (with no additional cognitive processing whatsoever) is conscious perceptual processing.

Recurrent Processing Theory

A sensory theory of conscious processing, RPT attributes perceptual awareness to first order sensory processing. This means, according to RPT, that following the initial feedforward sweep of sensory activity, recurrent processing in the same occipito-temporal areas *is* conscious perceptual processing (Lamme 2015, 2010, 2006, 2004, 2003, Lamme et al. 2000). Put differently, an integrated percept (contained in occipito-temporal areas, constructed from recurrent processing loops) is the necessary and sufficient condition for perceptual experience. In the eyes of RPT, the true NCC therefore should be found roughly in the 100-200ms time window, following the initial feedforward sensory sweep. Proponents of RPT view all processing following this 100-200ms time window as post-perceptual processing - this of course includes the late 300-600ms time window that GNWT claims houses the true NCC. Further, RPT claims that recurrent processing is not necessarily sufficient for conscious access, but is sufficient for perceptual experience. Therefore, proponents of RPT endorse the existence of non-accessible ‘phenomenal’ consciousness - which is said to be “pure” experience (Lamme 2004, Block 1995). This experience is thought of as simply perceptual experience with no metacognition whatsoever; in other words, perceptual experience with no higher-level awareness of such perceptual experience. The idea of phenomenal consciousness, in the eyes of a GNWT supporter, does not exist. It makes no sense, with respect to GNWT, to describe the existence of a “nonaccessible” perceptual experience (Naccache 2018). The idea of ‘phenomenal consciousness’ is just preconscious processing to GNWT - information is salient enough to be processed, but the fronto-parietal network is not involved (so no perceptual awareness) (Lamme 2006). Proponents of RPT and GNWT agree that late, widespread fronto-parietal activity indexed by the P3 ERP is necessary for accessing perceptual content. The disagreement, again, regards the existence of perceptual experience that is experienced, but not accessed by any other cognitive systems (that become involved later in time). In the case of conscious vision, RPT would claim that GNWT wants to explain how people *know* they see, rather than explaining *how* people see.

Integrated Information Theory

IIT (Tononi et al. 2016, Koch et al. 2016, Koch 2004, Crick & Koch 1995) makes similar predictions to RPT regarding the timing and neuroanatomical regions involved in consciousness. For proponents of IIT, integrated information is responsible for conscious content. Therefore, phi (ϕ), the amount of integrated information in a system is a measure of how much conscious content is in a system. Phi refers to the maximum amount of cause-effect power; the processes with the highest phi values are important for conscious processing. The theory outlines the existence of multiple ‘complexes’ (systems of integrated information) such that the processes underlying perceptual experience are carried out by the complex with the highest phi value. Relatedly, the theory supports the existence of a purely “phenomenal” type of conscious processing, different from “access” conscious processing - which is said to be needed for higher-level awareness and report. IIT is not limited to brains or biological materials; one could measure the (nonzero) phi-value of a highly complex computer. Such a computer could not function like today’s (most powerful) computing systems; but a conscious computing system would necessarily require the information processing complexity found in biological materials. Therefore, proponents of IIT claim that today’s computers lack the information-processing capabilities to induce conscious processing; however, future computers that are orders of magnitude more powerful than computers today (and designed to have cause-effect power on themselves) can be conscious. Additionally, IIT does not make any hard predictions about where in the brain conscious processing arises, but it currently postulates a “posterior cortical hotzone” involving occipital-temporal-parietal cortex as a likely region of interest due to its grid-like architecture (Koch 2018). However, according to IIT, the major complex (responsible for processing perceptual experience) could move around to various areas of cortical tissue. Importantly, this theory of conscious processing was constructed first by an analysis of the intrinsic nature of experience, resulting in four fundamental axioms: intrinsic existence (one's experience is one's own), composition (there is some structure), information (there is specificity), integration (the pieces are tied together), and exclusion (there is stuff not included, experience is bounded). Then, the theory was applied to brain science; as opposed to the other way around - using the brain science to then construct a theory (with axioms). For instance, IIT assumes the existence of 'qualia', a concept that refers to the specific characteristics that define each perceptual experience - which are necessarily unique (Koch et al. 2016).

Techniques for Comparing Conscious and Nonconscious Processing

Experimental approaches to isolating conscious perceptual processes utilize a variety of strategies for comparing conscious and nonconscious neural processing. Manipulations of perception alter crucial characteristics of stimuli such that the stimuli can be consciously perceived or not, depending on such characteristics. Examples of these perceptual manipulations include pattern masking (Dehaene et al. 2001), binocular suppression, binocular rivalry (Blake et al. 2002), binocular fusion, motion induced blindness (Bonneh 2010), and continuous flash suppression (Tsuchiya & Koch 2005). In all of these cases, it is the sensory processing of the stimuli that is disrupted, either by other non-critical stimuli (masks, competing stimuli in the other eye, moving vs. static stimuli, etc.) or by degrading the critical stimuli to threshold levels (i.e. in which the "signal" is barely distinguishable from the "noise").

Attentional manipulations can also cause changes in conscious *content* (change blindness and inattentional blindness) (Rensink et al. 1997, Chapman et al. 1964). For example, inattentional blindness and the attentional blink are two of the most popular attention-altering paradigms. Inattentional blindness is induced when attentional resources are allocated toward certain information such that other information (for which attentional resources are not allocated) is not processed deeply – resulting in its exclusion from perceptual experience. In other words, inattentional blindness is a paradigm that utilizes the fact that the attentional scope is finite – if some information is within the spotlight of attention, then there must be information *not* within the spotlight of attention. The information outside of the spotlight is not perceived. Similarly, the attentional blink paradigm utilizes the finite scope of temporal attention, inducing perceptual awareness of some stimuli, but not others. The attentional blink uses rapid serial visual presentation (RSVP) in which most of the images are distractor images, while some images are critical stimuli (which must, at least, be detected by participants). When two critical stimuli are presented in close temporal proximity (~200-600 ms apart), perception of the second critical stimulus is drastically impaired (Raymond et al. 1992). In contrast, when two critical stimuli are presented in the absence of such temporal proximity (600+ ms apart), perception of both critical stimuli remains intact.

In the vast majority of experiments that aim to isolate signals related to perceptual awareness, the contrastive approach is used. This approach contrasts neurophysiological data correlated with perceptual awareness, with neurophysiological data correlated with the lack of perceptual awareness. When using these paradigms that aim to isolate neural signatures of perception, subjects' self-reports are traditionally responsible for providing the behavioral data (and consequently contaminating the neural data). In other words, the general limitation of these paradigms in studying conscious perception is the fact that people must provide some verification of perceptual awareness (or lack thereof) - meaning that neural activity associated with perception *as well as* neural activity associated with deciding/reporting about perception is recorded when using these paradigms in a traditional manner. Not insignificantly, people are likewise prone to be imperfect when reporting upon perception, just like any other task. Thus, it seems thoroughly important to develop methods of dissociating verifiable perceptual awareness from the (possibly accurate) report of this perceptual awareness.

This limitation of traditional methods of isolating neurophysiological data confounds perception per se with some extra decision-making (about that perception). The solution to this issue is to decouple the task (subjects providing responses) from the manipulation of awareness. Put differently, does the neural difference that exists between seen and unseen stimuli *when people are completing a task on those stimuli* also exist between seen and unseen stimuli *when people are NOT completing a task on those stimuli*? This is precisely the question that the current study aims to answer. By orthogonally manipulating awareness (using visual pattern masking) and task (by having people respond to critical stimuli or noncritical stimuli), the study is aimed at dissociating perceptual processes from task-related processes (of working memory, language, motor preparation, etc.). When subjects provide responses about target stimuli, these stimuli are said to be task-relevant; and when subjects provide responses about nontarget stimuli, the target stimuli are considered to be task-irrelevant. Put differently, when completing a task on noncritical stimuli, the critical stimuli do not matter for task performance (deeming them task-irrelevant); whereas, when completing a task on critical stimuli, the critical stimuli matter greatly for task performance (deeming them task-relevant).

EEG and ERPs

Electroencephalography (EEG) is an electrophysiological methodology used in cognitive neuroscience. Scalp EEG is a noninvasive technique that measures changes in electrical potential (voltage) across the scalp - these changes in voltage correspond to activity in populations of cortical neurons. EEG is a fantastic method for examining signals with precise timing, as the temporal resolution of the technique is robust. On the other hand, the spatial resolution of scalp EEG is relatively poor; and thus, finding the neuroanatomical origin of a signal is often impossible. The event-related potential (ERP) is the product of a certain (common) type of EEG analysis. Time-locking the EEG signal to an observable event (stimulus onset or a button-press response) and then averaging the related electrophysiological signals across similar trials (or blocks, conditions, depending on experimental design) allows us to measure electrical brain activity related to these events, i.e. ERPs (Luck 2014, Hillyard & Kutas 1983).

P3b ERP

The late, occipito-temporal positivity, that correlates closely with task-relevant perceptual awareness, is called the P3b (or P300, P3, late-positive [LP]) event-related potential (ERP). This ERP has been widely studied for over fifty years, first being isolated by Chapman et al. (1964), when asking participants to complete a task on two visually-presented digits: “respond with which digit is smaller” (Chapman et al. 1964). Subsequently, this late positive signal was quickly distinguished into two broad distributions: the P3a and the P3b. The P3a was elicited when participants perceived an oddball stimulus while not actively attending; differently, the P3b was elicited when participants perceived an oddball stimulus while actively attending. The P3a ERP has a shorter latency and more frontal distribution than the larger amplitude, longer latency P3b ERP, which has an occipito-temporal distribution (Squires et al. 1975). More recently, the pair of electrophysiological signals have been given more nuanced descriptions, with evidence for the signals coding for difficulty of stimulus discrimination, rather than stimulus novelty (Kleih et al. 2011, Polich 2007, Comerchero et al. 1999). As the P3b ERP has been associated with top-down attentional and memory processes, it has been offered as a candidate neural correlate of consciousness, compared to the P3a ERP – more associated with bottom-up attentional processes (and not as strongly proposed as an NCC) (Batterink et al. 2012, Koivisto & Revonsuo 2010, Fisch et al. 2009, Gaillard et al 2009, Del Cul et al. 2007, Sergent et al. 2005).

No report paradigms

First Experiment

In recent years, the SCALP Lab Group (based at Reed College) has been involved with the development of no report paradigms that aim to isolate neural data associated with perceptual awareness. In 2014, Pitts et al. used an inattentional blindness paradigm to manipulate awareness of critical stimuli, while also manipulating task requirements for participants, across phases of the experiment. As shown in Fig 1, Pitts et al. used a sustained inattentional blindness paradigm that included two main types of stimuli: a ring of red dots, and a grid of small white lines. In the first phase of this experiment, participants were instructed to allocate their attention toward the ring of red dots, responding when one of the disks randomly became dimmer than the others (this occurred on 10% of trials). While engaged in this task, the small white lines in the center of the red ring ('twitchy lines' - in this phase, noncritical stimuli) had three possible configurations: random (no shape form), diamond, or square pattern. Immediately following Phase 1, participants completed an awareness questionnaire that probed participants' awareness of the noncritical ('twitchy line') stimuli. If participants were not aware of any shape patterns, the questionnaire served as a cue for participants - increasing the likelihood of noticing the noncritical stimuli in Phase 2 of this experiment. During this second phase, participants completed the exact same task as in Phase 1: respond when one of the red dots becomes dimmer than the others; the difference between these first two phases is the presumed awareness of noncritical stimuli in Phase 2, compared to Phase 1. Before moving on to Phase 3, participants recompleted the awareness questionnaire that was given in between Phase 1 and Phase 2. Then, in Phase 3, participants were asked to allocate attention toward the grid of 'twitchy lines', responding when the diamond pattern was presented, ignoring the ring of red dots (Pitts et al 2014 [Neuroimage]).

Thus, awareness and task-relevance were manipulated across these three experimental phases. In Phase 1, the red rings were task-relevant, while the 'twitchy lines' were task-irrelevant (i.e. not needed to complete the task on the red ring). Phase 2 had the same task-relevance architecture as Phase 1, but we infer awareness of the task-irrelevant ‘twitchy lines’ about which participants were prompted (in the awareness questionnaire). In Phase 3, the red ring became task-irrelevant and the ‘twitchy lines’ became task-relevant as participants were instructed to detect diamond patterns in the ‘twitchy line’ array. Importantly, the neural data associated with the square pattern (in the ‘twitchy lines’) remained untouched by (motor) response artifacts across all three phases as the square pattern was never responded to. Therefore, comparisons across awareness and task-relevance could take place. Half of participants were aware of the shape stimuli during Phase 1 (deemed the “noticer” group, as opposed to the “inattentionally blind” group), while all of the participants became aware of the shape stimuli during Phase 2. Pitts et al. found that the NCC in question, the P3 ERP, was *only* present in Phase 3, when the square stimuli became task-relevant (Pitts et al 2014 [Neuroimage]).

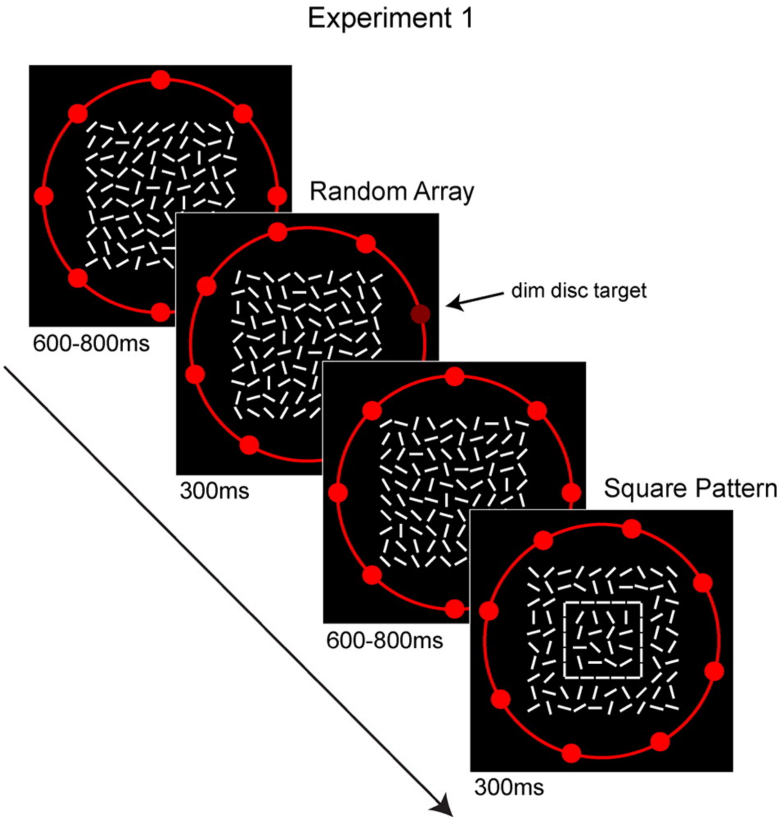

Fig 1. Pitts et al. (2014a) stimuli

Stimuli from (Pitts et al. 2014 [Neuroimage]) are shown. There were two stimulus types used: (1) the outer, red ring, and (2) the inner array of ‘twitchy lines’.

Second Experiment

Pitts et al. (2014 [Frontiers in Psychology]) then completed a second experiment, where instead of including a ring of red disks, the experimenters included red-colored patches of 'twitchy lines' (shown in Fig 2). All participants were shown colored patch and rectangle pattern stimuli before the experiment to prevent inattentional blindness to task-irrelevant stimuli. During the colored patch task, half of participants responded to the presence of 3 colored patches, and the other half of participants responded to the presence of 4 colored patches. Similarly, during the shape task, half of participants responded to vertical rectangle patterns, while the other half of participants responded to horizontal rectangle patterns. An important detail, the colored patches only appeared at locations that corresponded with rectangle edges during shape trials. Between completing the shape task and color patch task (the order of these tasks counterbalanced across participants), a similar awareness questionnaire was given. Only two participants were inattentionally blind to the rectangle patterns (during the color patch task), so they were disqualified from analyses. Again, the researchers found that the P3 ERP was present when stimuli were task-relevant, and absent when stimuli were task-irrelevant, despite all stimuli presumably being consciously perceived. This relationship between task-relevant and task-irrelevant stimuli is the same for color patch and shape pattern stimuli.

Pitts et al. decided to complete a third experiment. In this third rendition, the 'twitchy lines' had three types of arrangements: random, rectangle pattern, and ellipse pattern. Before completing any of the experimental procedure, the participants viewed all four possible shape stimuli: vertical and horizontal rectangles, and vertical and horizontal ellipses. For the entire experiment, half of participants attended to the rectangles and half of participants attended to the ellipses. Within these rectangle and ellipse groups, half of participants attended to the horizontal version of the shape, while the other half attended to the vertical version of the shape. Thus, stimuli can be categorized by similarity to target stimulus. Stimulus types in order of decreasing similarity: 'same shape different orientation', 'different shape same orientation', 'different shape different orientation'. The P3 showed a graded response, determined by target similarity; this, again, supports the notion that the P3 ERP is more related to task-relevance (i.e. post-perceptual processing) than to strictly perceptual processing. If the P3 component was related to perception per se, then it would be present during perceptual awareness and absent during lack of perceptual awareness.

Pitts et al. (2014) completed another study in which perceptual awareness and task-relevance were manipulated. Instead of using a sustained inattentional blindness paradigm, where perception is altered by manipulating attentional processes, the researchers elected to use a backward masking paradigm. The experimenters presented the same shape and color stimuli from previous experiments, but this time followed the critical stimuli by random arrays of ‘twitchy lines’ stimuli in order to disrupt sensory processing, altering conscious perception of the critical stimulus. While recording neural activity via EEG, the subjects were shown five different stimulus types: random ‘twitchy lines’ arrays, squares, diamonds, three red lines, or four red lines. The stimuli were either shown for 16ms or 300ms, each timing corresponding to a different level of perceptual awareness: unaware and aware, respectively. The manipulation of task was also straightforward. On shape task blocks, participants pressed a button when a diamond was detected; and during color task blocks, participants responded with button press when three red lines were detected. Participants completed two blocks of each task, alternating between the two; this was counterbalanced across participants. The researchers found, again, that the P3b ERP is not a consistent marker of the contrast between perceptual awareness and perceptual unawareness. Instead, they found a P3b component in some unaware, task-relevant data as well as the lack of a P3b component in some aware, task-irrelevant neural activity. Therefore, it seems as though the P3b is more associated with post-perceptual processes (related to task demands) than with perception per se.

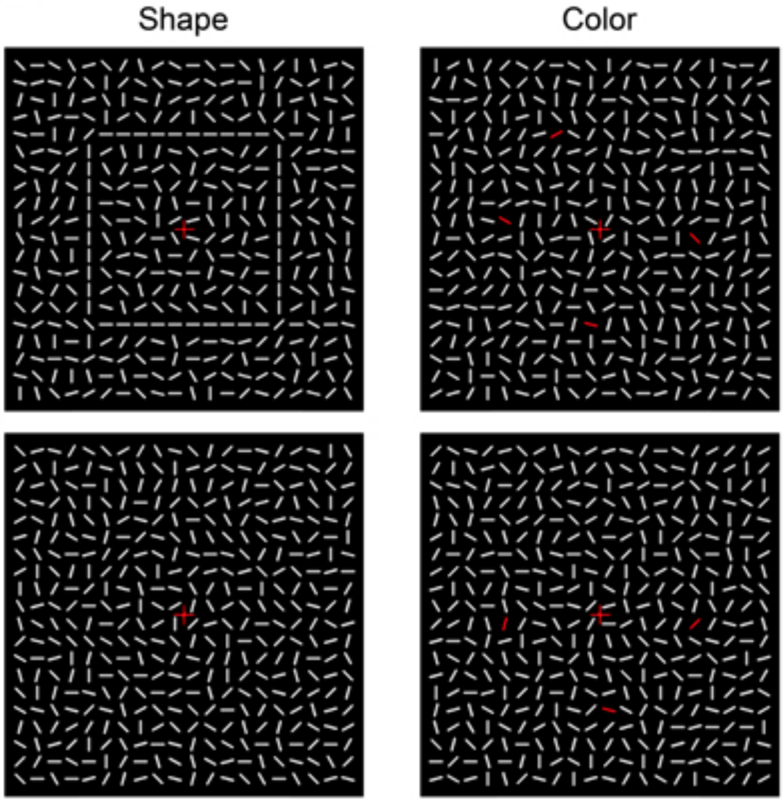

Fig 2. Pitts et al. (2014b) stimuli

Stimuli from Pitts et al. (2014 [Frontiers in Psychology]) are shown. There were two stimulus types used in the array of 'twitchy lines': (1) noncolored, closed shapes, and (2) red-colored patches of 'twitchy lines'.

Introduction to Current Study

The proposed experiment will critically inform our current understanding of visual perceptual awareness and put several leading empirical theories of consciousness to the test (Dehaene 2014, Dehaene et al. 2006, Dehaene et al. 2004, Dehaene et al. 2001, Dehaene et al 1998, Boly et al. 2017, Koch et al. 2016, Lamme 2006, Lamme 2003, Lau et al. 2011, Prinz 2012). Almost all previous studies in this area of research have compared neural activity elicited by physically identical stimuli that are consciously perceived on some trials and unconsciously processed on other trials (Dehaene et al. 2011, Bor et al. 2012). In such experiments, subjects are tasked with providing trial-by-trial reports of what they saw, e.g., by pressing buttons to indicate “I saw a face” or “I saw a house” or “I didn’t see anything”. Unfortunately, the main limitation of these report paradigms is that the brain recordings include the neural activity of interest, plus the neural activity associated with the subject’s reporting on what they perceived and it’s difficult to parse the two (Aru et al. 2012). No-report paradigms in consciousness research are scarce, however, this proposal seeks to directly change that (Frassle et al. 2014, Wilke et al. 2009). No-report paradigms are uniquely beneficial for studying the differences between conscious and unconscious processing (Tsuchiya et al. 2015). This is simply because no-report paradigms allow experimenters to usurp the main limitation of report paradigms. In the recent past, the SCALP Lab at Reed has developed several no-report paradigms, and the current project seeks to continue this trend on the path to isolating the neural correlates of perceptual awareness (Pitts et al. 2017, Pitts et al. 2015, Pitts et al 2014 [Frontiers in Psychology], Pitts et al 2014 [Neuroimage], Pitts et al. 2012).

The current study aims to identify consistent neural markers of conscious visual perception regardless of the task-relevance and the way in which awareness is manipulated (disruption of sensory processing vs. disruption of attention). The current study takes a seminal masking study as its base (Dehaene et al. 2001), first replicating its manipulation of perceptual awareness, then expanding the experiment to include a task-irrelevant condition in which the same (awareness) contrast is isolated from the task. The current experiments improve upon prior no-report research from Pitts et al. (2014

[Neuroimage], [Frontiers in Psychology]). First, these experiments utilize a standard paradigm – pattern masking, as opposed to inattentional blindness – with which it is easier to modulate perceptual awareness. Additionally, the current investigation instructs participants to do simpler, less-demanding tasks than previously used – this change minimizes the chance that NCCs are being altered by task-difficulty. Altogether, the current investigation is a more focused inquiry into the P3b's relation to perceptual awareness and task-relevance.

Conceptually, this is a straightforward manipulation: if the P3b ERP is the difference between aware and unaware perceptual processing when people are responding about their perception (of critical stimuli), will the P3b ERP also be the difference between aware and unaware perceptual processing when people are not responding about their perception (of critical stimuli)? Subjects completed two main visual tasks: a detection task in task-irrelevant blocks and a discrimination task in task-relevant blocks. Of course, during the task-irrelevant condition, participants were not responding about the critical stimuli. Instead, participants completed a trivially easy task where they pressed a button upon detecting a colored stimulus (appearing on ~10-20% of trials). Importantly, during the task-irrelevant condition, it is necessary to know that participants are indeed viewing the critical stimuli (that are not necessary for the color task). Thus, after completing the task-irrelevant condition, participants are queued on their memory of critical stimuli; since asking participants during the condition would defeat the whole purpose of a task-irrelevant condition. Receiving a delayed response about perception (that is encoded in some form of memory) is an improvement upon Pitts et al.'s previous work where either trial-by-trial responses accounted for subjects' responses, or subjects' completed a questionnaire regarding perceptual awareness. This questionnaire, again, does not allow the researchers to determine how often stimuli were perceived; therefore task-irrelevant stimuli could be perceived one time, or *every* time - asking participants about their perception *after-the-fact* does not allow for accurate estimates of frequency of perception.

Predictions

According to cognitive theories such as the GNWT, there are four distinct neural correlates of conscious perception: (1) activation of fronto-parietal circuits, (2) increased long-range neural synchrony, (3) a late P3b wave, and (4) late, high-frequency gamma-band oscillations (Dehaene 2014). Contrastingly, sensory theories (e.g. RPT, IIT) posit that the NCC according to cognitive theories is actually more associated with post-perceptual processes, compared with perceptual processes per se. So, these theories claim that the true NCC is localized within a "posterior cortical hot zone" - this is near the regions where low and mid-level visual processing are shown to occur (Koch et al. 2016, Boly et al. 2017). Of main interest here is the possible NCC, the P3b ERP. Cognitive and sensory theories make different predictions regarding the P3b ERP. Namely, GNWT claims that the P3b will be present during perceptual awareness, and absent (or largely diminished) during a lack of perceptual awareness. On the other hand, RPT claims that the P3b does not follow the behavior outlined by GNWT, but that recurrent processing in sensory areas does (meaning that this is the actually NCC). Further, sensory theories and cognitive theories of conscious processing make different predictions about the timing of the NCC, as shown in Fig 3. The current experiments were designed to test if the P3b (a late positive component) does indeed behave as a neural marker of conscious visual perception. Therefore, any result of these experiments informs the current understanding of visual perceptual awareness.

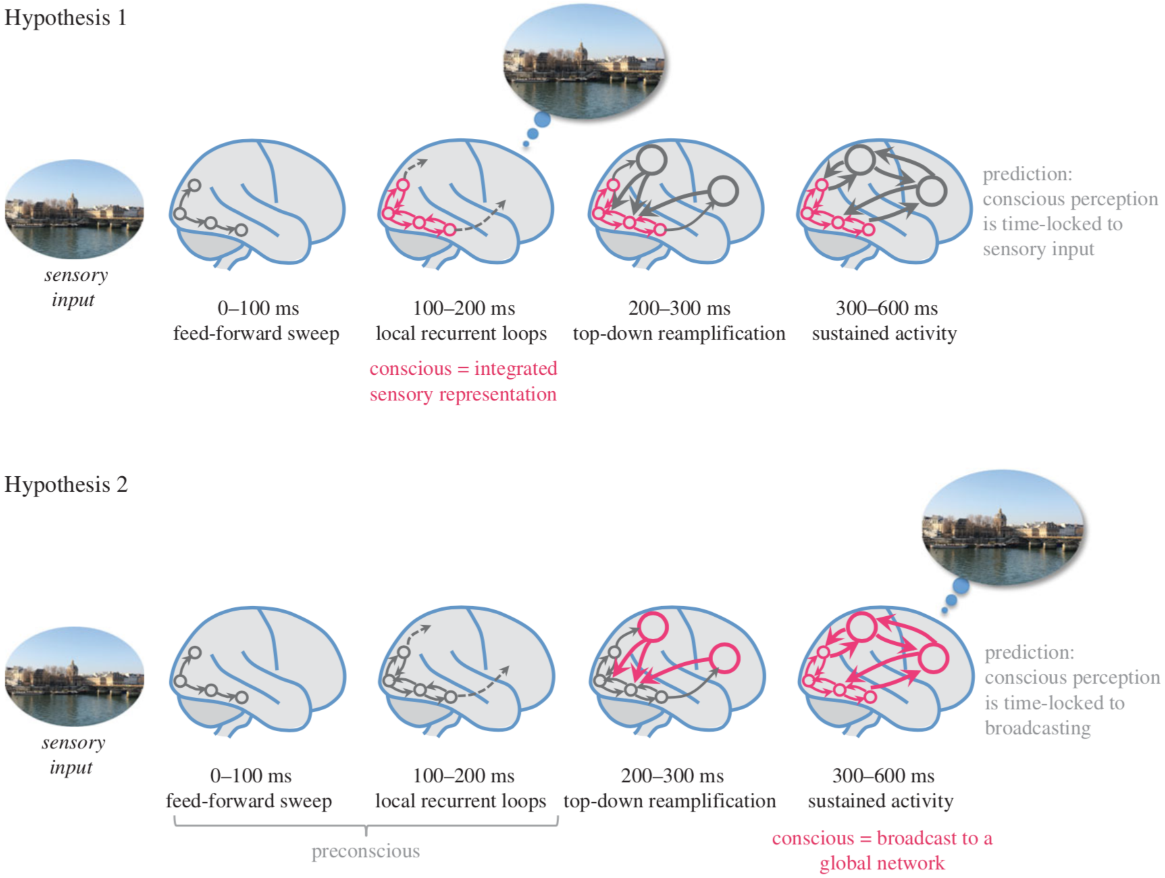

Fig 3. Different NCC Predictions

From Sergent (2018). One simple way to dissociate NCC predictions from different theories is by their proposed latency. Sensory theories, like RPT (top row; Hypothesis 1), predict that the NCC is located in recurrent loops in sensory areas, occurring roughly in the 100-200 ms time range. On the other hand, cognitive theories, like GNWT (bottom row; Hypothesis 2), predict that the NCC is reliant on widespread cortical processing, roughly occurring in the late 300-600 ms time range.

Methods – Experiment One & Two

Design Overview of Experiments

The current research is comprised of two nearly identical masking experiments. In both experiments the critical stimuli are line drawings of animals and objects, and these stimuli are rendered “seen” or “unseen” in different trials based on the temporal sequence of masks (scrambled line stimuli) and blanks surrounding the critical stimuli. The first experiment utilizes a red oval, while the second experiment utilizes a green circle as noncritical stimuli. There are two conditions in each experiment: (1) the task-irrelevant condition, where participants complete a detection task on the noncritical stimulus; and (2) the task-relevant condition, where participants complete a discrimination task on the critical stimuli. In each condition, the *exact* same stimuli are presented, and the only change that occurs across conditions is the task participants are instructed to accomplish. The second experiment’s design was based on a limitation of the first one: the noncritical stimuli in the first experiment were temporally and spatially related to the critical stimuli. Thus, in experiment one, the critical stimuli, during the task-irrelevant condition, may not have been quite task-irrelevant enough to properly test the main hypotheses. Therefore, in the second experiment the critical and noncritical stimuli were decoupled, both temporally and spatially – in order to render the critical stimuli completely irrelevant to the task in the task-irrelevant condition. There are two main comparisons of interest from these experiments: (1) the contrast between ‘seen’ and ‘unseen’ task-relevant critical stimuli, and (2) the contrast between ‘seen’ and ‘unseen’ task-irrelevant critical stimuli. The main question, in brief, is whether the neural difference (i.e. the P3b ERP) in the task-relevant comparison is also reflected in the task-irrelevant comparison.

Stimulus Presentation Paradigm

A sandwich masking paradigm, based on Dehaene et al. (2001), was used in both experiments. This straightforward manipulation of stimulus order reliably modulates perceptual visibility of the target stimulus, according to previous experiments as well as preliminary pilot testing in the SCALP Lab. ‘Seen’ trials were those where the critical stimulus was lightly masked - meaning it was flanked by blank images. Conversely, ‘unseen’ trials were those where the critical stimulus was heavily masked - meaning it was flanked by mask images. The same mask image was used for the entirety of both experiments, and was constructed by using a mixture of line segments from critical stimulus images, overlaid upon each other. Blank trials were also included so the neural activity associated with perception of *only* the critical stimulus could be isolated. To accomplish this, the neural activity associated with a blank stream (mask-blank-blank-blank-mask, for instance) was subtracted from the neural activity associated with a critical stimulus stream (mask-blank-stimulus-blank-mask, for instance) - leaving *just* the neural activity associated with the critical stimulus. In total, there were four trial types: (1) ‘seen’ critical stimulus, (2) ‘seen’ blank, (3) ‘unseen’ critical stimulus, and (4) ‘unseen’ blank (shown in Fig 4). All of these trial types appeared with 25% frequency; and the noncritical colored stimulus was presented on 10-20% of trials (evenly distributed across trial types).

In Experiment One, a response screen consisting of a single question mark (on a blank white background) was presented subsequent to the main 5-image visual stream. During the task-irrelevant condition, the response screen lasted for 900 ms (participants only provided responses on about 20% of trials); and during the task-relevant condition, the response screen lasted until the participant provided a button-press response. In Experiment Two, a blank response screen consisting of only the (constant) red fixation dot was presented for 1000 ms (1 s). This change to Experiment Two’s response screen is to minimize timing (and procedural) differences between the task-irrelevant and task-relevant conditions.

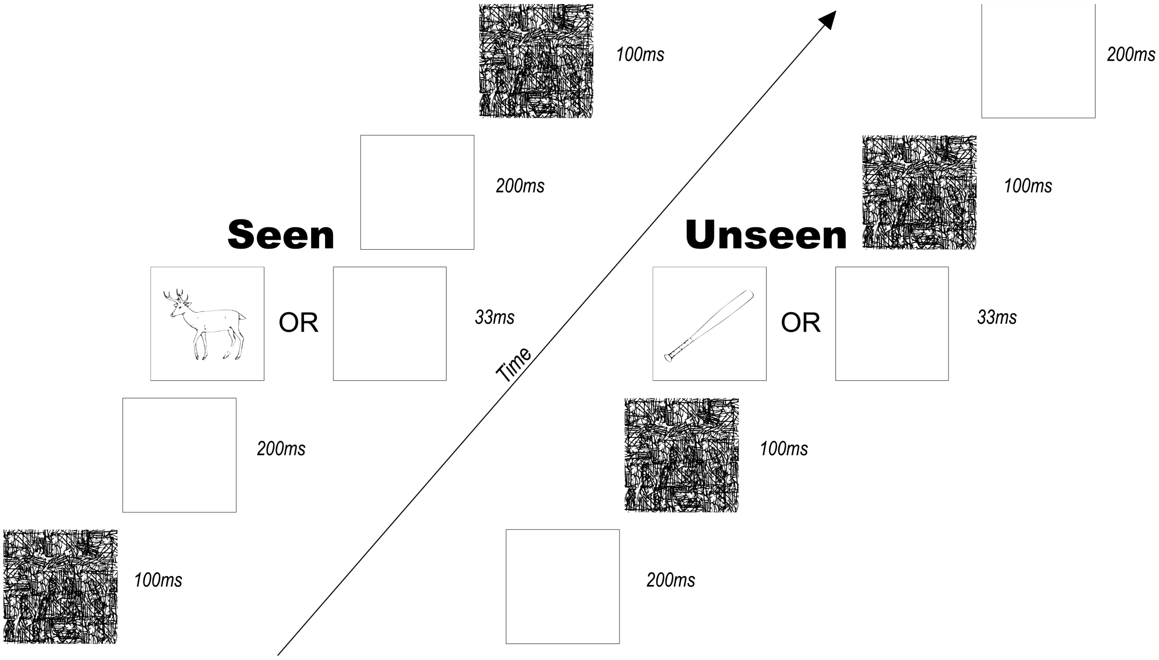

Fig 4. Stimulus Masking Procedure

Structure of trials across both experiments, using sandwich masking to manipulate perceptual awareness. Half of all trials were presented as ‘seen’ trials (left side of figure), while the other half of trials were presented as ‘unseen’ trials (right side of figure). In half of all ‘seen’ and ‘unseen’ trials, instead of presenting a critical stimulus, a blank screen was presented (to allow subtraction of mask-elicited brain activity for ERP analysis). Stimulus timing does not change across trial types, only the order of stimuli is altered - this results in the manipulation of visual perceptual awareness.

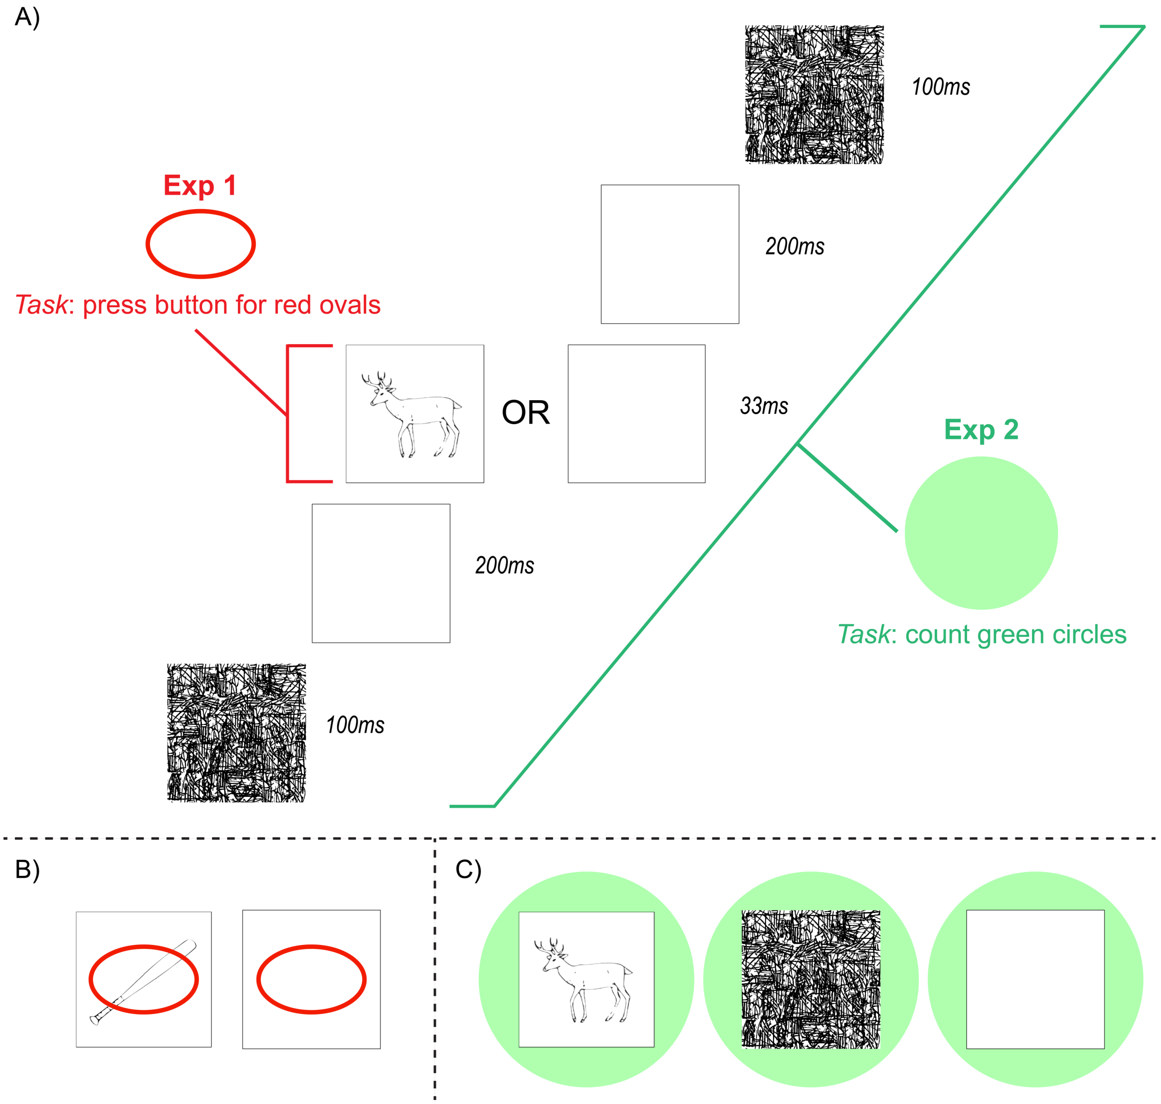

Fig 5. Noncritical stimuli

This figure shows examples of colored stimuli across both experiments. Participants responded on these colored stimuli during the task-irrelevant condition (in each experiment). *(Panel A)* In Experiment One (left), the colored noncritical stimulus was a red oval that appeared superimposed spatially and temporally with the critical stimulus (or blank). Contrastingly, in Experiment Two (right), the colored noncritical stimulus was a large green circle that appeared in the background; the green circle could appear during presentation of the critical stimulus, the masks, the blanks, or anywhere else in the mask-stimulus-mask sequence or subsequent inter-trial interval period. *Panel B* displays the possible locations of the noncritical stimulus in Experiment One. *Panel C* displays the possible locations of the noncritical stimulus in Experiment Two.

Stimuli

The critical stimulus set contained a total of 32 images (16 animals, 16 objects). Each participant was presented all 32 critical stimuli. Critical stimuli were divided into 4 groups (A, B, C, and D), each group consisting of 4 animals and 4 objects. Each participant, according to their subject number, was assigned a stimulus group (A, B, C, or D) for each trial type: (1) seen, task-relevant; (2) unseen, task-relevant; (3) seen, task-irrelevant; and (4) unseen, task-irrelevant. Therefore, in one condition (either task-relevant or task-irrelevant), exactly half of the stimuli were presented to the participant. And, given perfect manipulation of awareness via masking, exactly half of the presented stimuli were seen by the participant in each condition. This, again, is because half of the stimuli were presented using light masking ('seen' trials) and half of the stimuli were presented using heavy masking ('unseen' trials). The *mask* images were constructed using line segments from the *critical stimuli* images, overlaid upon each other such that no obvious shapes could be perceived in the image, providing no object-based perceptual information to participants.

All stimuli used in the experiment were images, 625x625 pixels (px), presented on top of a white background (BenQ 120Hz monitor, 1920x1080 px). Stimulus presentation was carried out by Psychophysics Toolbox Version 3 for MATLAB (Kleiner et al. 2007), and participants were seated roughly 70 cm from the monitor.

Incidental Memory Test (IMT)

Like previously described, lightly masked stimuli comprised the ‘seen’ trials, whereas heavily masked stimuli comprised the ‘unseen’ trials. The main purpose of the IMT was to verify a proper ‘seen’ versus ‘unseen’ comparison - which relies on the ‘seen’ trials actually being seen, and the ‘unseen’ trials actually being unseen. In the task-irrelevant condition, since button-press responses were not collected for critical stimuli, the IMT ensures that stimuli presented on ‘seen’ trials were indeed seen, and vice versa for the stimuli presented on ‘unseen’ trials. The IMT consisted of presenting 24 images to each participant, and for every image, a ‘yes’ or ‘no’ response was collected via button press. A ‘yes’ response indicated the participant remembered seeing that *exact* image during the previous condition; and a ‘no’ response indicated the participant did not remember seeing that *exact* image during the previous condition (made up of several blocks). Of the 24 images shown to participants, 8 were the stimuli from the ‘seen’ trials, 8 were the stimuli from ‘unseen’ trials, and 8 were never shown in any part of the experiment (i.e. foils). Perfect performance on the IMT meant that participants correctly accepted all 8 ‘seen’ stimuli, and correctly rejected all 16 ‘unseen’ and foil stimuli. In sum, the IMT allowed the neural comparisons of interest to be validated, given participants’ accurate task performance.

Animal/Object Task

During the task-relevant condition, participants carried out the animal, object, nothing task. In this task, participants were briefly presented the same stream of five images, after which they were cued about whether they detected an animal, object, or nothing. Subjects responded using a computer keyboard, specifically, the left, down, and right arrow keys - each key corresponding to one of the three response choices. The critical stimuli, in this condition, are now task-relevant. Subjects’ button-press responses are key-consistent across all blocks of the animal, object, nothing task; and the button-press responses are counterbalanced across subjects in order to account for any artifacts due to sensory-motor mapping of the button-ordering.

EEG Recording and ERP Pre-Processing

All data was gathered using equipment and laboratory space owned by the Reed College Psychology department. Brain electrical activity was recorded non-invasively from the scalp using a custom equidistant 64-channel “Easycap”. Electrode signals were sampled at a 500 Hz digitization rate and amplified by BrainVision “Professional BrainAmp” amplifiers, with a high cutoff filter at 150 Hz (data was later low-pass filtered offline at 30 Hz), using BrainVision Recorder. During task performance, eye movements were monitored by vertical and horizontal EOG recordings. An average of the mastoid electrodes served as the reference for all scalp channels during recording. Each session lasted 140–180 min, including electrode cap preparation and recording time. All trials that contained a noncritical stimulus presentation were excluded from all analyses, as the focus was on neural responses to critical stimuli.

Using BrainVision Analyzer 2, ERPs were time-locked to the onset of critical stimuli (or blank stimuli, in place of critical stimuli) and baseline corrected from −200 to 0 ms. The left and right horizontal EOG channels were re-referenced as a bipolar pair. Trials containing eye movements, blinks, or other muscle artifacts in a –600 to +600 ms interval surrounding stimulus-onset were discarded from further analysis. Artifact detection was accomplished semi-automatically via per-subject adjustment of the following peak-to-peak thresholds: eye movements (50 µV, 50 ms steps, in bipolar HEOG), blinks (100 µV, 200 ms steps, in VEOG and FP1), and muscle noise (150 µV, 200 ms steps, all remaining channels). On average, 15.8% of trials were rejected due to a combination of these artifacts. Any participant with greater than 30% of trials rejected per stimulus type per phase was excluded from analysis due to an insufficient number of trials to obtain reliable ERPs (zero subjects were excluded). Individual electrodes showing extended periods of noise in the raw EEG were removed and replaced by interpolated signals from surrounding channels using topographic spherical splines.

Mass Univariate Analysis

For data-driven statistical analysis of EEG data, Mass Univariate Analysis Toolbox, with MATLAB, was used (Groppe et al. 2011). Difference waves obtained by subtracting ‘unseen’ trials (after blank/mask subtraction) from ‘seen’ trials (also after blank/mask subtraction) were subjected to repeated measures, two-tailed t-tests for each condition (task-irrelevant and task-relevant) using the cluster-based permutation analysis for multiple comparisons at the 5% level. The ERP data was down-sampled to 100 Hz before these mass univariate analyses. All time points between 100 and 600 ms for all 59 scalp electrodes (excluding HEOG and VEOG) were included in the analyses (2950 total comparisons per mass univariate analysis). The mass univariate tests indicated the presence or absence of the hypothesized ERP differences, which were used to select time windows for ERP analysis.

Experiment One

Procedure

Subjects

14 college-age adults (18-24 years) with normal visual acuity and color vision volunteered from the Reed College community. Participants provided informed consent to participate prior to completing any of the experimental procedure. The participants completed two tasks and an incidental memory test (IMT) while neural activity was recorded with scalp electroencephalography (EEG). At the end of the experimental procedure, participants were debriefed about the procedure and the purpose of the IMT. Additionally, all participants were compensated with $20. No participants were excluded due to excessive EEG artifacts (or any other reason). All procedures were approved by the Reed College Institutional Review Board.

Color Task

Participants completed 4 blocks (232 trials each block) of the oval detection task. In this task, participants were briefly presented a series of five images (totalling 633 ms), after which they were cued for a brief interval (900 ms) asking if they detected a colored oval. Half (n=7) of the participants completed a version of the task with just red ovals, where the possible responses were a button press (“yes”, an oval was presented) or no button press (“no”, an oval was not presented). The other half (n=7) of the participants completed the oval detection task with two possible colored ovals: one red and one green. These subjects had three possible responses in this case: one button press for detecting a red oval, a different button press for detecting a green oval, and no button press for detecting no oval. The two groups of participants did not show different behavioral or neural data, therefore the groups were collapsed in data analysis and the two tasks treated as exemplars of a more general ‘color detection’ task. The purpose of this colored oval task was to provide subjects with an easy task that was unrelated to the presentation of the critical stimuli, in order to ensure that subjects were actually paying attention to the (critical) stimuli presented. If there was no task at all in the task-irrelevant (i.e. passive viewing), then we would run the risk of subjects engaging in a range of task-unrelated activity (e.g. thinking about lunch, about how they will wash their hair after the experiment, covertly naming the animal/object pictures), that would impact neurophysiological data and potentially ruin the task manipulation.

Incidental Memory Test (IMT)

Following the task-irrelevant condition, participants completed the IMT.

Animal/Object Task

Following the IMT, participants completed 4 blocks (232 trials each block) of the animal, object, nothing task.

Results

Behavioral Results

There were two types of behavioral data collected across the experimental procedure: (1) task performance (color task, animal/object task), and (2) IMT performance (after task-irrelevant condition). Perfect color task performance includes correctly withholding a button-press response when noncritical stimuli were absent (i.e. non-colored trials) and correctly responding with a button-press when critical stimuli were presented. Perfect animal/object task performance includes correctly detecting animals, objects, and blanks when presented during 'seen' trials. Perfect IMT performance means selecting images presented during 'seen' trials, and rejecting images presented during 'unseen' trials and images never presented (foils). All behavioral results are displayed below.

Stimulus Type	**Average Accuracy (n=14)**
Lightly masked	90.2%
Heavily masked	6.2%
Catch	4.5%

Table 1. Exp 1 IMT Results

This table displays the post-task-irrelevant incidental memory test data from (n=14) participants, separated by stimulus type. The manipulation of awareness, from these data, seems to be functioning reliably.

Color task average accuracy (n=14)	96.6%

Table 2. Exp 1 Color Task Results

This table displays the task-irrelevant condition's color task performance.

Stimulus Type	**Average Accuracy (n=14)**
All lightly masked stimuli	91.66%
All heavily masked stimuli	52.02%
All blanks	92.26%
Lightly masked animals	90.80%
Lightly masked objects	91.29%

Table 3. Exp 1 Animal/Object Task Results

This table displays the results from the task-relevant condition's animal/object task, separated by stimulus type. Based off these data, the manipulation of awareness seems to be functioning reliably.

Event-Related Potential (ERP) Results

All results include neural data associated with the critical stimuli (or blanks in place of critical stimuli, on ‘blank’ trials), and exclude all neural data from trials where a colored stimulus was presented. All ERP figures depict grand-averaged data from (n=14) participants. All participants completed every part of the experimental procedure, thus all comparisons are within-subjects. The first comparison of interest was between ‘seen’ and ‘unseen’ task-relevant trials (Fig 6). The second main comparison of interest was between ‘seen’ and ‘unseen’ task-irrelevant trials (Fig 7). For both of these main comparisons, four time windows (based on MUA results) are shown for topographic distributions. Specifically, four 40 ms time windows, centered on the canonical P3b latency, map the progression of scalp electrical activity that differentiates seen vs. unseen trials. Next, the ERP analyses focused on the ‘seen’ versus ‘unseen’ task-irrelevant comparison – but this time, in the front of the head. To this end, we see a frontal negativity, centered roughly around 300 ms (Fig 8). Finally, we compared ‘unseen’ task-relevant trials with ‘unseen’ task-irrelevant trials in order to assess differential brain activity related to the task (even for unconsciously processed stimuli), and found a late positive difference for ‘unseen’ task-relevant trials (Fig 9).

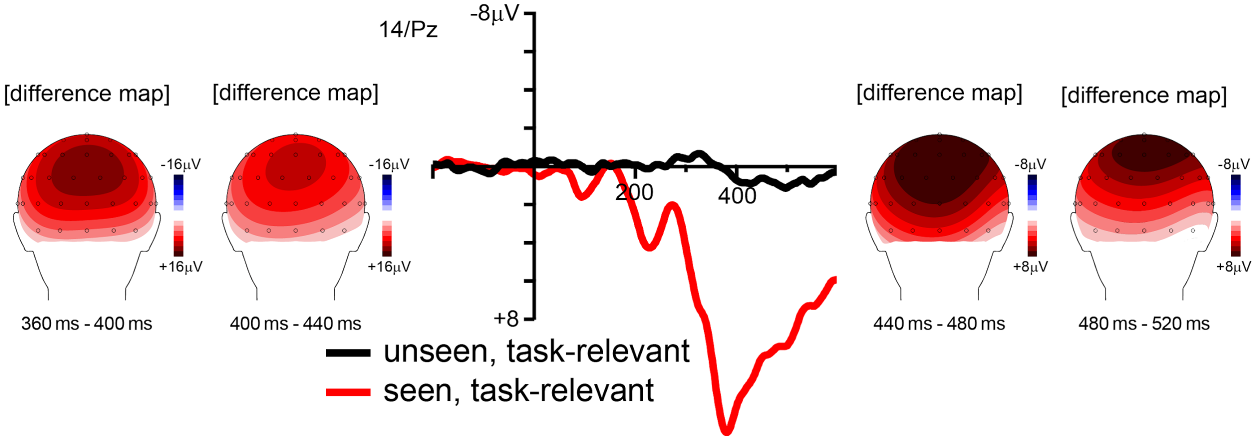

Fig 6. Exp 1 ERP: task-relevant ‘seen’ vs. ‘unseen’ – P3b
Data from (n=14) grand-averaged participants at central parietal electrode location 14/Pz. This is the comparison between ‘seen’ task-relevant trials and ‘unseen’ task-relevant trials. Scalp maps display ‘seen’ minus ‘unseen’ ERP data at four successive time windows (each averaged over 40 ms intervals). All four scalp maps display topographies resembling the canonical occipito-parietal P3b distribution.

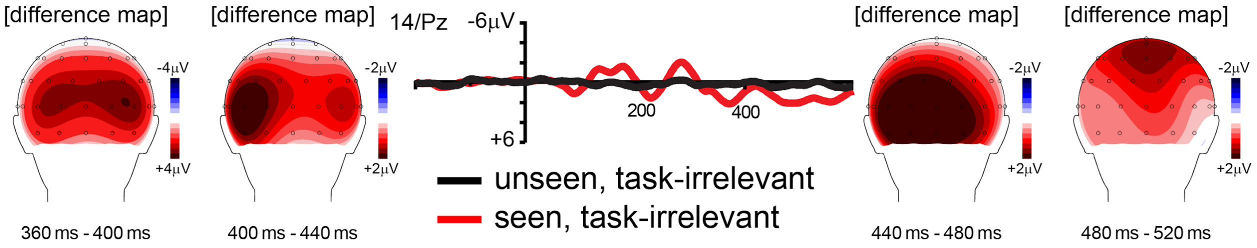

Fig 7. Exp 1 ERP: task-irrelevant ‘seen’ vs. ‘unseen’ – late posterior positivity
Data from (n=14) grand-averaged participants at central parietal electrode location 14/Pz. This is the comparison between ‘seen’ task-irrelevant trials and ‘unseen’ task-irrelevant trials. Scalp maps display ‘seen’ minus ‘unseen’ ERP data at four successive time windows (each averaged over 40 ms intervals). The final scalp map (480 ms – 520 ms) roughly resembles a typical P3b distribution, while the previous three do not (but see Fig 8 & discussion).

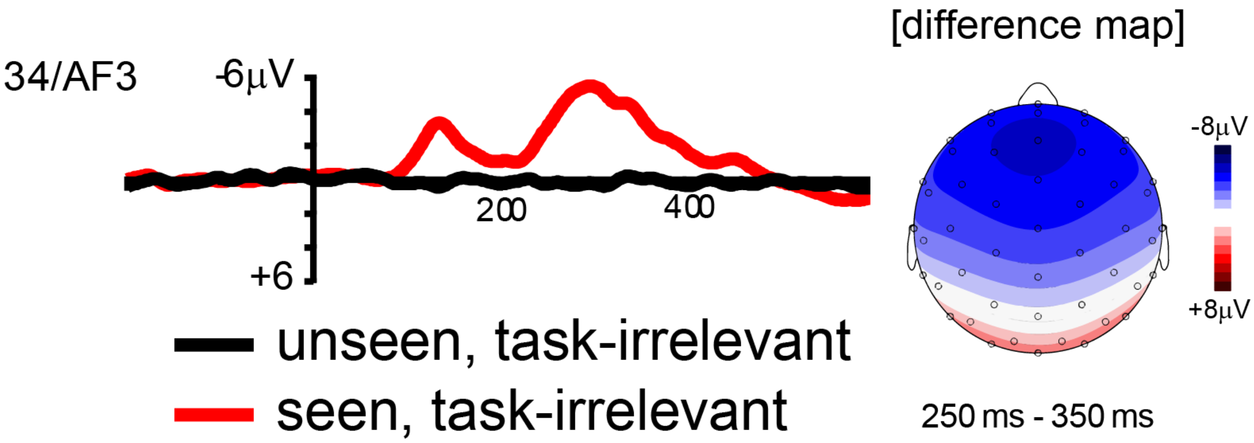

Fig 8. Exp 1 ERP: task irrelevant ‘seen’ vs. ‘unseen’ – frontal negativity
Data from (n=14) grand-averaged participants at frontal electrode location 34/AF3. This is the comparison between ‘seen’ task-irrelevant trials and ‘unseen’ task-irrelevant trials, this time, taken at the front of the head. The scalp map displays ‘seen’ minus ‘unseen’ ERP data averaged over a 100 ms time window, centered around the second negative deflection in the ERP to ‘seen’ task-irrelevant critical stimuli (which was statistically significant according to the MUA; see Fig 11).

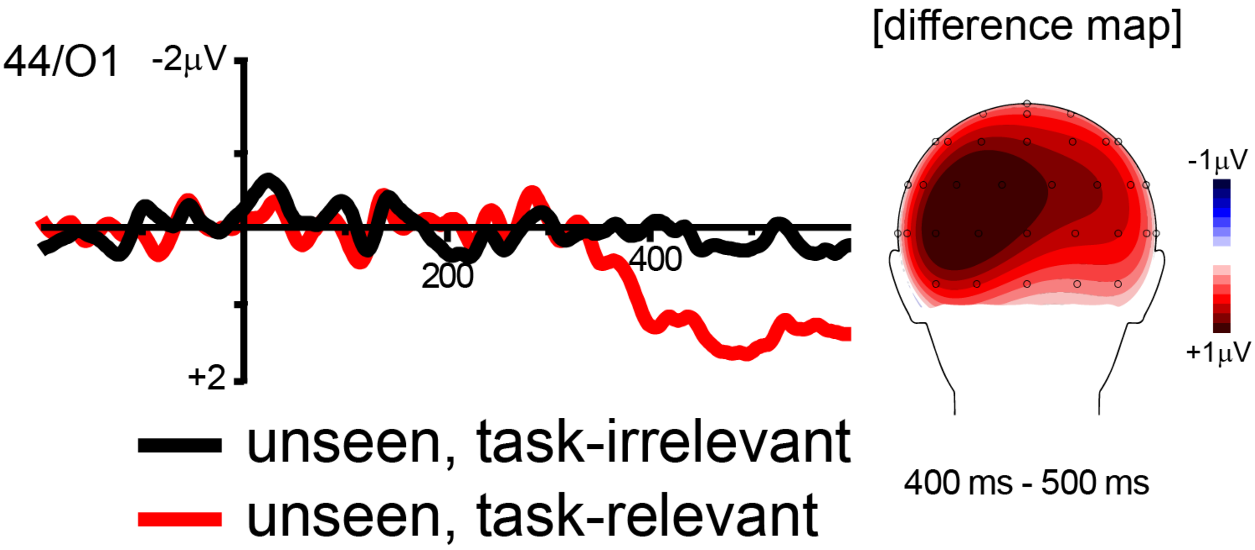

Fig 9. Exp 1 ERP: task-relevant ‘unseen’ vs. task-irrelevant ‘unseen’
Data from (n=14) grand-averaged participants at occipital electrode location 44/O1. This is the comparison between ‘unseen’ task-relevant trials and ‘unseen’ task-irrelevant trials. Scalp map displays an averaged 100 ms time window, showing ‘unseen’ task-relevant minus ‘unseen’ task-irrelevant data.

Mass Univariate Analysis (MUA) Results

The following plots display the results of mass univariate cluster analysis. MUA is a data-driven way to determine significant electrical amplitudes in difference waves of interest. Here, these are the differences between ‘seen’ and ‘unseen’ task-relevant trials (Fig 10) and ‘seen’ and ‘unseen’ task-irrelevant trials (Fig 11). For each individual participant, difference waves are computed for all electrode locations, which were entered into the Mass Univariate ERP toolbox for MATLAB. The program tests for consistent amplitude differences at every time point and every electrode, correcting for multiple comparisons. Again, this analysis was comprised of 2950 total comparisons. All colored cells on the plot correspond to statistically significant neural activity during a certain time window, at a specific electrode location. The figures are divided into three main horizontal bands: one corresponding to the left hemisphere, midline, and the right hemisphere.

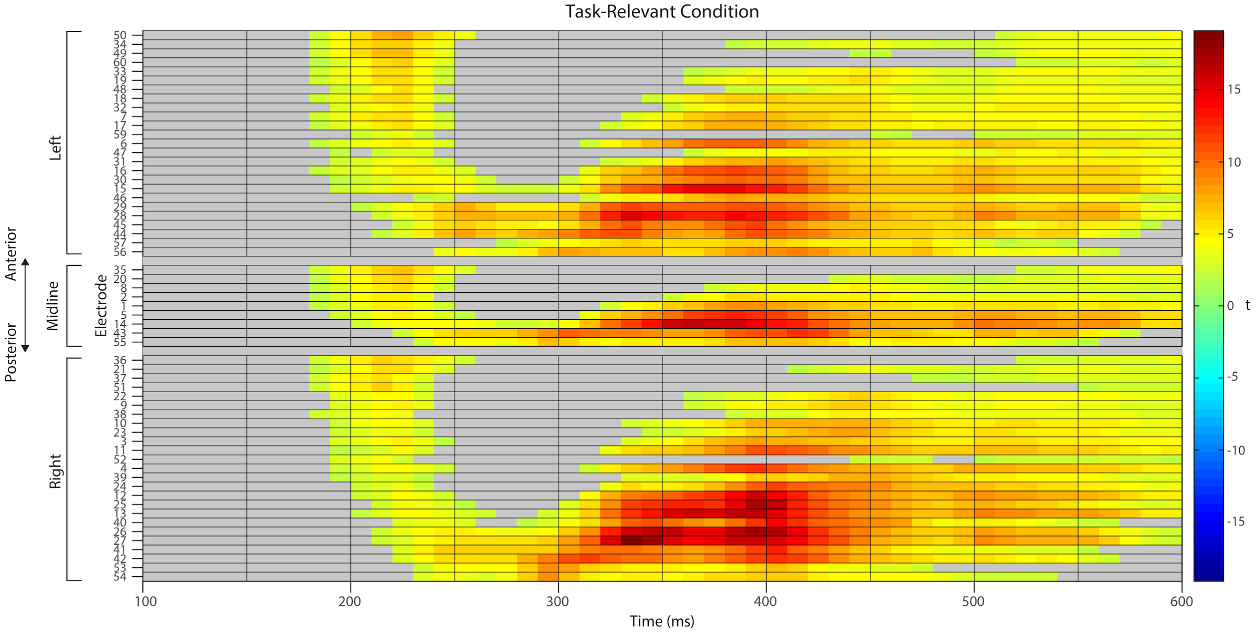

Fig 10. Exp 1 MUA: task-relevant ‘seen’ minus ‘unseen’

Mass univariate plot displaying statistically significant (p > .05) time points across critical time window (100-600ms post-critical stimulus onset). All neural data associated with ‘unseen’ task-relevant trials was subtracted from that of the ‘seen’ task-relevant trials, resulting in a plot of the significant differences between ‘seen’ and ‘unseen’ task-relevant perceptual processes.

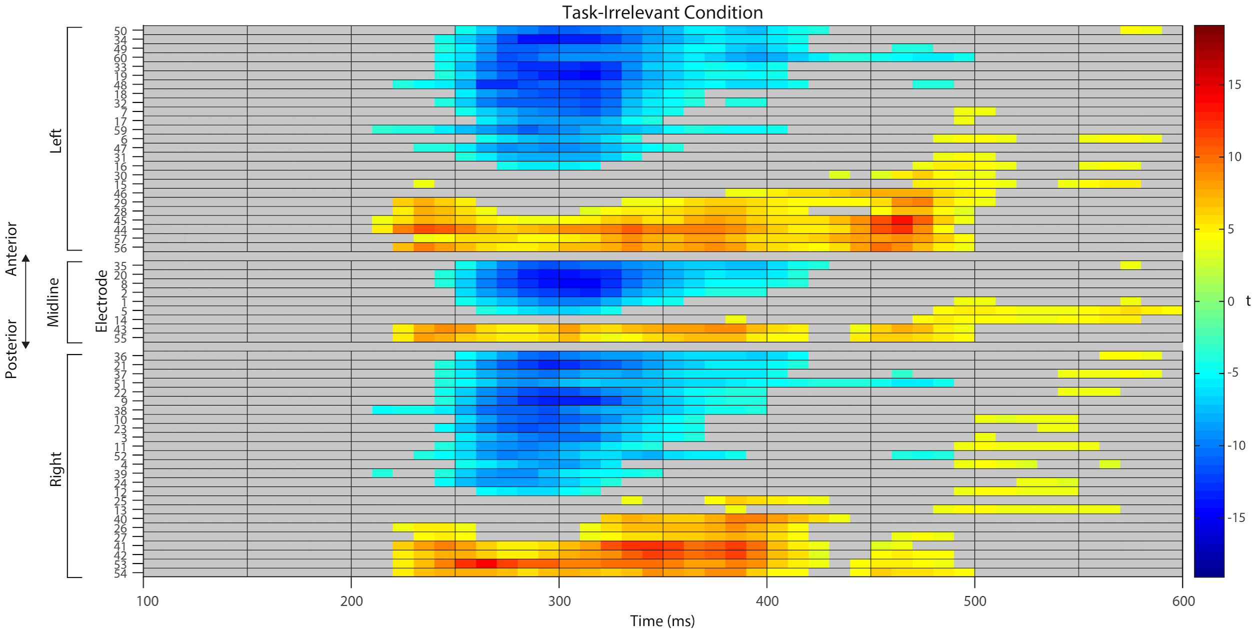

Fig 11. Exp 1 MUA: task-irrelevant 'seen' minus 'unseen'

Mass univariate plot displaying statistically significant (p > .05) time points across critical time window (100-600ms post-critical stimulus onset). All neural data associated with 'unseen' task-irrelevant trials was subtracted from that of the 'seen' task-irrelevant trials, resulting in a plot of the significant differences between 'seen' and 'unseen' task-irrelevant perceptual processes.

Results Summary

In brief, the results replicated previous findings in the task-relevant condition: a P3b was present during 'seen' task-relevant trials, and absent during 'unseen' task-relevant trials. So, when stimuli are task-relevant, the P3b ERP correlates with perceptual awareness. For the task-irrelevant condition, the P3b was massively reduced, and barely present, if at all. The scalp distribution of the P3b in the task irrelevant condition was not consistent with the P3b in the task relevant condition. However, a unique frontal negativity was observed in the task irrelevant condition, which appears to have overlapped the P3b in time, potentially skewing its scalp distribution.

Experiment Two

Design Overview and Rationale

The purpose of this second experiment is to account for a limitation of the first experiment. Simply put, critical stimuli, when seen but task-irrelevant, may not have been *just* passively viewed, but actively inhibited. The neural data alludes to participants making some decision *not* to respond when critical stimuli were presented and irrelevant to the task. Specifically, the frontal negativity observed in the seen task-irrelevant neural data has a similar latency and distribution to a 'no-go' negativity, a well-known ERP associated with withholding a response (in other words, deciding to the make the response of 'none') (Folstein et al. 2007, Breton et al. 1988). In a simple (go/no-go) task where one either responds or does not respond, neural activity associated with the 'no-go' (non)response is more negative than the neural activity associated with 'go' response (Gajewski et al. 2013). In short, this is a limitation of the previous experimental design where noncritical stimuli, when they appeared, did so during the critical stimulus presentation. This design resulted in a spatial and temporal overlap between noncritical (task-relevant at the time) stimuli and critical (task-irrelevant at the time) stimuli. Temporal and spatial overlap in the presented image may have led to some neural processing overlap of these different stimuli. This overlap between stimulus types is a problem mostly because it changes the neural activity associated with perceiving the critical stimuli. It is possible that 'seen' task-irrelevant stimuli elicit a P3b response, yet such a response was not detectable in the current data because the frontal negativity cancelled (and/or distorted) the occipito-parietal late positive response (P3b). Therefore, from the first experiment, we cannot confidently conclude that task-irrelevant critical stimuli were in fact task-irrelevant. Thus, despite the large amplitude reduction of the P3b in the task-irrelevant condition, there may have still been a statistically reliable (and much more modest) P3b in this condition, therefore, no strong claim about the P3b as a marker of perceptual awareness can be made based on Experiment One alone.

In the second experiment, we designed the stimulus presentation stream such that noncritical stimuli were completely temporally decoupled from the critical stimuli. In this experiment, the noncritical (large, green circle) stimuli appear on about 10% of trials, and can randomly appear at any temporal location within a trial. Additionally, the large, green circle was bigger than the normal mask-stimulus-mask stream, and presented in the background - minimizing spatial overlap. These changes in timing and spatial layout, again, differ from the previous experiment where critical and noncritical stimuli were presented too similarly to be truly irrelevant. Finally, instead of tasking subjects with pressing a button upon detection of the noncritical stimuli (in the task irrelevant condition), which may lead to more response inhibition when the noncritical stimulus is absent, we introduced a counting task instead. The goal of the second experiment was to make the noncritical stimuli as unconnected as possible to critical stimuli, in order to prevent the "no-go" inhibitory activity that may have disrupted our measures of the P3b in experiment one. Achieving this goal would make critical stimuli, during the task-irrelevant condition, truly task-irrelevant.

Subjects

3 college-age adults (18-24 years) with normal visual acuity and color vision volunteered from the Reed College community. Participants provided informed consent to participate prior to completing any of the experimental procedure. The participants completed two tasks and two incidental memory tests (IMT) (one after the task-irrelevant condition and one after the task-relevant condition) while neural activity was recorded with scalp electroencephalography (EEG). At the end of the experimental procedure, participants were debriefed about the procedure and the purpose of the IMT. Additionally, all participants were compensated with $20. No participants were excluded due to excessive EEG artifacts (or any other reason). All procedures were approved by the Reed College Institutional Review Board.

Color Task

This task is virtually identical to Task 1 from Experiment One, but with the following key changes. Instead of a hollow colored oval, the color change appears via a solidly colored green circle. Additionally, instead of only being able to appear during the 33ms that critical stimuli (or blanks) are presented, the green color change was able to occur at any (randomly chosen) point during the trial. Also, instead of occupying a space within the 625x625px that the masks, blanks, and critical stimuli inhabited, the green circle was bigger than the mask-stimulus-mask stream area (and presented in the background for a duration of 33 ms). Finally, instead of requiring an immediate button-press response, subjects were tasked with counting the number of green circles during each block, and had to enter the total into the keyboard after each block ended. Thus, these color changes, in relation to those from Experiment One, are less relevant to the critical stimuli as they no longer are temporally and spatially linked, and did not require immediate responses. In other words, the green color change here is less task-relevant to the critical stimuli; and that is exactly the aim of this second experiment. Finally, instead of 4 blocks of 232 trials each, Experiment Two utilized 15 blocks of 64 trials each.

Trials, in this experiment, were comprised of the exact same 633ms 5-item stream as in Experiment One as well as a 1s blank screen. The blank screen was included because participants, in the task-relevant condition, provided a response for every trial, during this blank period. In order to make the most straightforward comparison, stimuli were the *exact* same across conditions, meaning that an unused response screen was included in the task-irrelevant condition. So, the total duration of an individual trial was 1633ms. On trials where there was a green color change (always 33 ms duration), it was equally likely to happen during the 633ms of mask-stimulus-mask stream as it was the 633ms immediately after this mask-stimulus-mask stream (the last 367ms of the blank response window was not available for green color changes). This distribution of possible color-change timing also ensured the lowest amount of task-relevance as the color change was able to happen with or without the mask-stimulus-mask stream - the two shared no causal relationship. Participants were instructed to count how many times a green circle appeared during the previous 64-trial block - possible answers were always 4, 6, and 8 - and participants responded with a single button press at the end of each block.

Incidental Memory Test

After completing the 15 blocks (64 trials each) of the colored oval task, participants were given the incidental memory test. Additionally, after completing the 15 blocks of the animal/object task, the participants were given the same surprise memory test (on different stimuli).

Animal/Object Task

Participants completed 15 blocks (64 trials each block) of the animal, object, nothing task between the first and second IMT.

Results

Preliminary Behavioral Results

There were two types of behavioral data collected across the experimental procedure: (1) task performance (color task, animal/object task), and (2) IMT performance (after both task-irrelevant and task-relevant conditions). Perfect color task performance includes correctly responding with button-press indicating the number of green circles presented during the experimental block. Perfect animal/object task performance includes correctly detecting animals, objects, and blanks when presented during 'seen' trials. Perfect IMT performance means selecting images presented during 'seen' trials, and rejecting images presented during 'unseen' trials and images never presented (foils). All behavioral results are displayed below in Table 2, except for Animal/Object accuracy, due to technical issues regarding the accuracy evaluation algorithm. Strangely, participants achieved higher accuracy on the task-irrelevant IMT, compared to the task-relevant IMT. This means that participants had better memory for stimuli they were *not* asked to provide responses for, compared to stimuli that they *were* asked to provide responses for. The current small sample size does not elucidate the origin of this result.

Stimulus Type	**Average Accuracy (n = 3)**
Lightly masked	75%
Heavily masked	16.67%
Catch	20.83%

Table 4. Exp 2 Task-irrelevant IMT Results
This table displays the post-task-irrelevant incidental memory test data from (n=3) participants, separated by stimulus type.

Stimulus Type	Average Accuracy (n = 3)
Lightly masked	70.83%
Heavily masked	20.83%
Catch	20.83%

Table 5. Exp 2 Task-relevant IMT Results

This table displays the post-task-relevant incidental memory test data from (n=3) participants, separated by stimulus type.

Task	Average Accuracy (n=3)
Color	88.9%
Animal/Object	n/a (see above)

Table 6. Exp 2 Behavioral Results

The behavioral performance for all tasks from Experiment 1 are displayed. The right column displays the average accuracy (for n=3 participants) in the task listed in the left column.

Preliminary ERP Results

All results include neural data associated with the critical stimuli (or blanks in place of critical stimuli, on ‘blank’ trials), and exclude all neural data from trials where a colored stimulus was presented. All ERP figures depict grand-averaged data from (n=3) participants. It is important to keep in mind that this is data from (n=3) participants, compared with (n=14) participants in Experiment One. All participants completed every part of the experimental procedure, thus all comparisons are within-subjects. Like in Experiment One, the main two neural contrasts of interest were between ‘seen’ and ‘unseen’ task-relevant processing (Fig 12), and between ‘seen’ and ‘unseen’ task-irrelevant processing (Fig 13). ERP results were computed using the same time windows as the ERP results from Experiment One: using a progression of four time windows (averaged over 40 ms time periods) to view the topographical changes over time. Also, the same comparison was made between ‘seen’ and ‘unseen’ task-irrelevant neural data, at the front of the head (Fig 14). Due to the smaller sample size in Experiment Two, the neural contrast between ‘unseen’ task-relevant and ‘unseen’ task-irrelevant trials was not explored; this comparison will be made when the signal-to-noise ratio is higher (due to more participants).

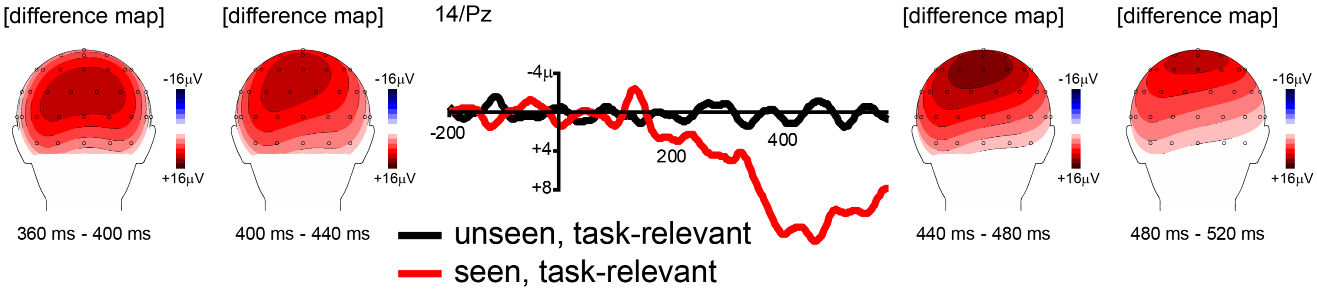

Fig 12. Exp 2 ERP: task-relevant ‘seen’ vs. ‘unseen’ – P3b
Data from (n=3) grand-averaged participants at central parietal electrode location 14/Pz. This is the comparison between ‘seen’ task-relevant trials and ‘unseen’ task-relevant trials. Scalp maps display ‘seen’ minus ‘unseen’ ERP data at four successive time windows (each averaged over 40 ms intervals). The three final scalp maps display topographies resembling the canonical occipito-parietal P3b distribution, and the first scalp map (360-400 ms) largely resembles such a distribution.

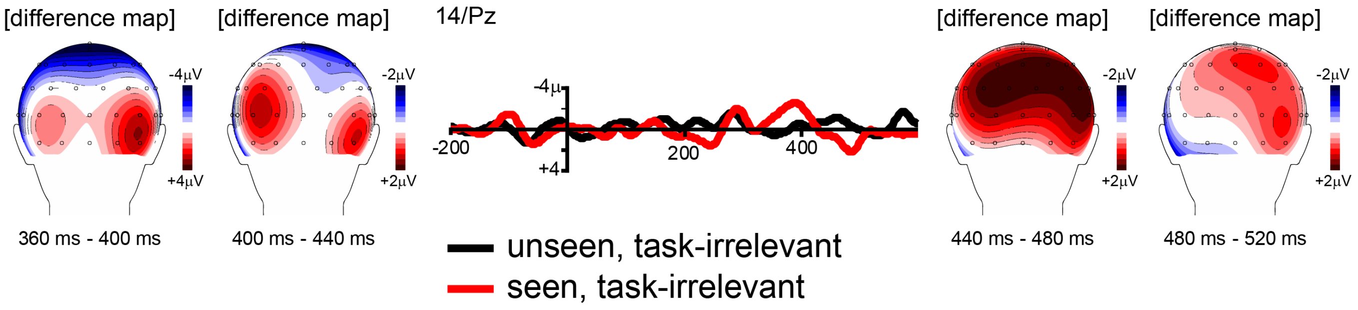

Fig 13. Exp 2 ERP: task-irrelevant ‘seen’ vs. ‘unseen’ – late posterior positivity
Data from (n=3) grand-averaged participants at central parietal electrode location 14/Pz. This is the comparison between ‘seen’ task-irrelevant trials and ‘unseen’ task-irrelevant trials. Scalp maps display ‘seen’ minus ‘unseen’ ERP data at four successive time windows (each averaged over 40 ms intervals). None of these four scalp maps resemble a P3b distribution (see Fig 6, 12). During the 440-480 ms time window, the distribution of the late positive difference slightly resembles the distribution of late positive differences in the task-irrelevant condition from Experiment One (Fig 7).

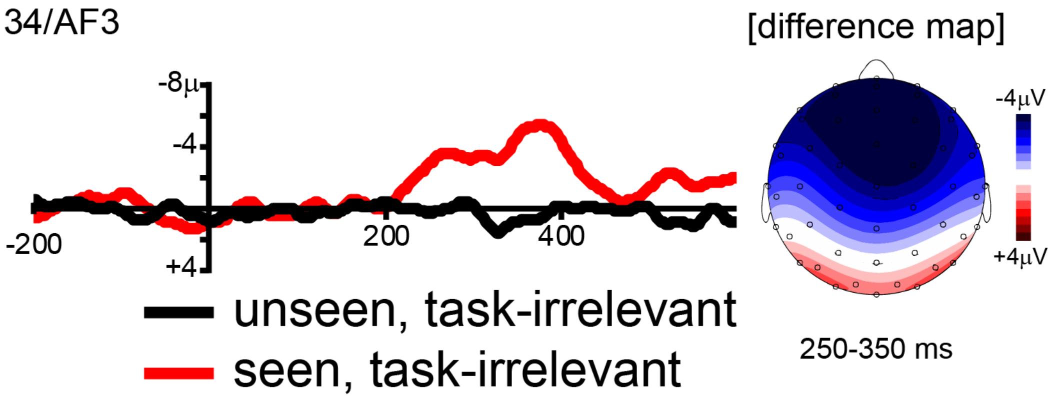

Fig 14. Exp 2 ERP: task irrelevant ‘seen’ vs. ‘unseen’ – frontal negativity
Data from (n=3) grand-averaged participants at frontal electrode location 34/AF3. This is the comparison between ‘seen’ task-irrelevant trials and ‘unseen’ task-irrelevant trials, this time, taken at the front of the head. The scalp map displays ‘seen’ minus ‘unseen’ ERP data averaged over a 100 ms time window, centered around the second negative deflection in the ERP to ‘seen’ task-irrelevant critical stimuli. This negative difference in Experiment Two is smaller than the same negative difference in Experiment One (Fig 8).

Discussion

Replicating Dehaene et al. (2001)

Experiment One

Dehaene et al. (2001) used a visual pattern masking paradigm that contrasted neural activity associated with visual perceptual awareness with neural activity associated with lack of visual perceptual awareness. Like intended, the current study replicated the results of Dehaene et al. (2001). When participants completed a task on critical stimuli (the task-relevant condition), the neural difference between visual perceptual awareness and lack thereof, is the P3b ERP. Put differently, the P3b ERP was present when participants viewed lightly masked ('seen') stimuli and absent when participants viewed heavily masked ('unseen'), during the task-relevant condition. In brief, this result is not surprising; but, replicating these past results is a crucial step toward validating results in the novel, task-irrelevant condition.

Experiment Two

From preliminary results, this (nearly identical) version of the masking experiment seems consistent with the task-relevant condition from Experiment One. Across the two task-relevant conditions, the only changes were slight ones to the stimuli: in Experiment Two, the noncritical stimuli became temporally and spatially decoupled from the critical stimuli; they also became more infrequent. Despite these changes, stimuli across experiments remained nearly identical, and similar task-relevant and task-irrelevant results were found.

The novel task-irrelevant condition

Experiment One

When critical stimuli were task-irrelevant, participants were asked to complete a task on the noncritical red ovals that appeared on ~20% of trials. Participants provided a button-press response when a trial contained a red oval, and withheld response on trials where no red oval was presented.

Does the P3b ERP describe the difference between task-irrelevant 'seen' and 'unseen' trials?

Unfortunately, it is difficult to tell if there is a P3b ERP while participants consciously perceived the critical stimuli but provided responses for the noncritical stimuli only. As depicted in Fig 7, the (posterior) neural difference between 'seen' and 'unseen' task-irrelevant stimuli is a positive one. However, this positive difference does not share the latency and distribution of a typical P3b ERP. Only in the 480-520 ms time window does the topography of the positivity resemble that of the central, occipito-parietal P3b. Contrastingly, in the task-relevant condition (Fig 6), all shown time windows display the canonical P3b distribution, showing that such a neural response is possible at a latency of ~360 ms. Importantly, the amplitude of neural responses was significantly decreased in the task-irrelevant condition, compared to the task-relevant condition (more than 8 times smaller).

So, it is possible that there is a diminished, delayed P3b in the task-irrelevant condition when participant view the lightly masked critical stimuli.

Frontal negativity in response to task-irrelevant critical stimuli

It is also possible that a P3b neural response was undetectable due to spatial and temporal overlap of the frontal "no-go" N2. The statistical results from the MUA (Fig 10 & Fig 11) show very similar patterns of positive differences in some electrodes during the

P3b time-frame, and all locations that differ, also happen to overlap with the tail-end of the frontal negativity that was uniquely observed in this the task irrelevant condition.

Experiment Two

From preliminary Experiment Two ERP results, it seems as though the methodological changes did not change the neural results from Experiment One. Unsurprisingly, the neural difference between ‘seen’ and ‘unseen’ task-relevant processes is the P3b ERP (Fig 12). However, again it is difficult to tell if the P3b ERP reflects the difference between ‘seen’ and ‘unseen’ task-relevant neural data. A slightly diminished frontal negativity was found in the task-irrelevant condition (compared to Experiment One), when comparing ‘seen’ and ‘unseen’ stimuli; therefore, it is possible that the current design changes worked in the direction that they were intended: reducing the relevance of the critical and noncritical stimuli such that there is less chance for inhibition of task-irrelevant critical stimuli. The frontal negativity was not extinguished, though, as was the goal of this second experiment. This preliminary finding complicates the interpretation of the task-irrelevant findings. Even with less task-relevant stimuli, there is still a frontal negative deflection that resembles that of a no-go inhibition negativity, i.e. a no-go N2 (Folstein et al. 2007, Breton et al. 1988).

Methodological challenges of no-report paradigms

As it turns out, measuring awareness in the absence of any extra cognitive processing is a difficult task. In other words, our goal is to correlate neural activity with perceptual awareness *per se*; and this means we want to isolate neural processes associated with simply perceiving information, without any memory encoding, and without providing any overt response (like a button-press). Why not, for instance, just present images to people and record their brain activity? Passive viewing (perceiving visual information, while providing no behavioral responses) is difficult to use for one main reason. How would we ensure that participants are indeed seeing what we want them to see? If passive viewing is used, there is no behavioral data with which to combine (and validate) the neural data. Participants could even be fixated on critical stimuli, yet are mind-wandering – which would alter the neural signature associated with perception of those stimuli. This is precisely the difficulty with designing no-report paradigms where the neural activity associated with some event is not relevant to the task participants are completing: we want to isolate the marker of *just* perceptual awareness, but we don't know what that marker is. Put differently, an inherent limitation of no-report paradigms is also their purpose: the fact that participants do not provide responses to the critical stimuli.

It seems that both the P3b and the N2 ERPs are consequences of conscious perceptual processing, and therefore not the NCC *per se*. For task-relevant information processing, the P3b ERP has been associated with a range cognitive processes from stimulus novelty, to discrimination difficulty, and to general cognitive task demands (Kleih et al. 2011, Polich 2007, Comerchero et al. 1999, Squires et al. 1975, Chapman et al. 1964). So far, there has not been research demonstrating that a P3b ERP can be elicited with a task-irrelevant perceptual processing. This could be due to the frontal "no-go" N2 ERP: which is associated with inhibiting behavioral responses to stimuli as well as to task-irrelevant perceptual processing (Folstein et al. 2007, Breton et al. 1988, Gajewski et al. 2013, Schendan et al 2012).

Future Directions

Is it possible to make the critical and noncritical stimuli even less related to each other? Decoupling these stimuli spatially and temporally (in Experiment Two) was not enough to eradicate the frontal negativity – indicative of cognition about task-irrelevant critical stimuli (Schendan et al. 2012). Cognition about task-irrelevant critical stimuli is precisely what we are trying to avoid by using a no-report paradigm. So, what would it take to separate these stimuli cognitively? Perhaps separate visual streams for critical and noncritical stimuli, or a higher working memory load during the task-irrelevant condition. Perhaps adding cognitive (memory) load during the task-irrelevant condition will minimize cognition about critical stimuli. Modulating working memory load was toyed with throughout the design process, as a previous working experimental design included a rotating disk that participants had to view and respond with a button-press when the rotating disk changed directions. This difficulty in isolating task-irrelevant perceptual processes may have ecological significance – perhaps task-irrelevant perceptual processes are in fact not common (or existent). In other words, maybe processes related to perceptual awareness are always task-relevant to some degree. Regardless of the truth of this claim, these results force one to ponder how rich our perceptual experience really is. It often seems as though we perceive irrelevant information – but is this information truly irrelevant to our perceptual experiences? Do we see “extra” stuff? Or do we see exactly the “right amount” of stuff? This question is an exciting one for the future of consciousness studies.

More immediately, the current investigation adds breadth to our understanding of such a ubiquitous ERP. The P3b may still be indicative of perceptual awareness per se, but this seems unlikely, as it is extensively linked to task-relevant cognitive processes (Pitts et al. 2014 [Neuroimage], Schendan et al 2012, Kleih et al. 2011, Polich 2007, Comerchero et al. 1999, Squires et al. 1975, Chapman et al. 1964). Therefore, the P3b might even be a better ERP for medical BCI technologies than previously thought. If the P3b is always task-relevant, then it is always indicative of task-relevant perceptual processes (which would strongly point to vegetative patients being alive and perceptually aware, compared to just alive), and would be easily usable in various electrophysiologically-based neurotechnologies. Further, a greater understanding of this ERP lends itself to a greater understanding of many neuropsychological conditions – for instance, schizophrenia. Understanding the task-relevant properties of ERPs like the P3b and the N2 could lend itself to a greater understanding of positive and negative schizophrenic symptoms, for instance, as the task-relevance of perceptual processes may be especially informative for these phenomena. In sum, the current investigation is one that informs a critical neurophysiological theory of conscious processing, medical practices that determine the path of a patient's life (and their loved ones), and it augments our view on the fundamental characteristics of our human experience.

References

Aru, J., Bachmann, T., Singer, W. & Melloni, L. Distilling the neural correlates of consciousness. *Neuroscience & Biobehavioral Reviews* **36**, 737–746 (2012).

Aru, J. & Bachmann, T. Still wanted-the mechanisms of consciousness! Frontiers in Psychology 6, (2015).

Aru, J. et al. Local Category-Specific Gamma Band Responses in the Visual Cortex Do Not Reflect Conscious Perception. Journal of Neuroscience 32, 14909–14914 (2012).

Baars, B. J. *A cognitive theory of consciousness*. (Cambridge University Press, 1995).

Baars, B. J. & Franklin, S. How conscious experience and working memory interact. *Trends in Cognitive Sciences* **7,**166–172 (2003).

Bar, M. & Ullman, S. Spatial Context in Recognition. *Perception* **25**, 343–352 (1996).

Bar, M. Visual objects in context. *Nature Reviews Neuroscience* **5**, 617–629 (2004).

Batterink, L., Karns, C. M. & Neville, H. Dissociable Mechanisms Supporting Awareness: The P300 and Gamma in a Linguistic Attentional Blink Task. *Cerebral Cortex* **22**, 2733–2744 (2012).

Blake, R. & Logothetis, N. K. Visual competition. *Nature Reviews Neuroscience* **3**, 13–21 (2002).

Block, N. On a confusion about a function of consciousness. *Behavioral and Brain Sciences,18*(02), 227 (1995).

Boly, M. *et al*. Are the Neural Correlates of Consciousness in the Front or in the Back of the Cerebral Cortex? Clinical and Neuroimaging Evidence. *The Journal of Neuroscience* **37**, 9603–9613 (2017).

Bonneh, Y. S. *et al.* Motion-induced blindness and microsaccades: Cause and effect. *Journal of Vision* **10**, 22–22 (2010).

Bor, D. & Seth, A. K. Consciousness and the Prefrontal Parietal Network: Insights from Attention, Working Memory, and Chunking. *Frontiers in Psychology* **3**, (2012).

Breton, F., Ritter, W., Simson, R. & Vaughan, H. G. The N2 component elicited by stimulus matches and multiple targets. *Biological Psychology* **27**, 23–44 (1988).

Chapman, R. M. & Bragdon, H. R. Evoked Responses to Numerical and Non-Numerical Visual Stimuli while Problem Solving. *Nature* **203**, 1155–1157 (1964).

Cohen, M. A. & Dennett, D. C. Consciousness cannot be separated from function. *Trends in Cognitive Sciences* **15,**358–364 (2011).

Cohen, M. A., Cavanagh, P., Chun, M. M. & Nakayama, K. The attentional requirements of consciousness. Trends in Cognitive Sciences 16, 411–417 (2012).

Comerchero, M. D. & Polich, J. P3a and P3b from typical auditory and visual stimuli. *Clinical Neurophysiology* **110**, 24–30 (1999).

Crick, F. & Koch, C. Are we aware of neural activity in primary visual cortex? *Nature* **375**, 121–123 (1995).

de Graaf, T. A., Hsieh, P.-J. & Sack, A. T. The 'correlates' in neural correlates of consciousness. Neuroscience & Biobehavioral Reviews 36, 191–197 (2012).

Dehaene, S., Changeux, J.-P., Naccache, L., Sackur, J. & Sergent, C. Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends in Cognitive Sciences* **10**, 204–211 (2006).

Dehaene, S. & Changeux, J.-P. Experimental and Theoretical Approaches to Conscious Processing. *Neuron* **70**, 200–227 (2011).

Dehaene S, Changeux P. Neural Mechanisms for Access Consciousness. In The Cognitive Neurosciences (ed M Gazzaniga). *MIT Press*, 1145–1153 (2004).

Dehaene, S., Kerszberg, M. & Changeux, J.-P. A neuronal model of a global workspace in effortful cognitive tasks. *Proceedings of the National Academy of Sciences* **95**, 14529–14534 (1998).

Dehaene, S. *et al*. Cerebral mechanisms of word masking and unconscious repetition priming. *Nature Neuroscience* **4**, 752–758 (2001).

Dehaene, S. Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition* **79**, 1–37 (2001).

Dehaene, S. *Consciousness and the brain: deciphering how the brain codes our thoughts*. (Viking, 2014).

Dehaene, S., Charles, L., King, J.-R. & Marti, S. Toward a computational theory of conscious processing. Current Opinion in Neurobiology 25, 76–84 (2014).

Dehaene, S., Lau, H. & Kouider, S. What is consciousness, and could machines have it? Science 358, 486–492 (2017).

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

Farooqui, A. A. & Manly, T. When attended and conscious perception deactivates fronto-parietal regions. Cortex (2017).

Fisch, L. *et al.* Neural "Ignition": Enhanced Activation Linked to Perceptual Awareness in Human Ventral Stream Visual Cortex. *Neuron* **64**, 562–574 (2009).

Folstein, J. R. & Van Petten, C. Influence of cognitive control and mismatch on the N2 component of the ERP: A review. *Psychophysiology*, (2007).

Frassle, S., Sommer, J., Jansen, A., Naber, M. & Einhauser, W. Binocular Rivalry: Frontal Activity Relates to Introspection and Action But Not to Perception. *Journal of Neuroscience* **34**, 1738–1747 (2014).

Gaillard, R. *et al.* Converging Intracranial Markers of Conscious Access. *PLoS Biology* **7**, e1000061 (2009).

Gajewski, P. D. & Falkenstein, M. Effects of task complexity on ERP components in Go/Nogo tasks. *International Journal of Psychophysiology* **87**, 273–278 (2013).

Groppe, D. M., Urbach, T. P. & Kutas, M. Mass univariate analysis of event-related brain potentials/fields I: A critical tutorial review: Mass univariate analysis of ERPs/ERFs I: Review. *Psychophysiology* **48**, 1711–1725 (2011).

Hillyard, S. A. & Kutas, M. Electrophysiology of Cognitive Processing. Annual Review of Psychology 34, 33–61 (1983).

İşcan, Z. & Nikulin, V. V. Steady state visual evoked potential (SSVEP) based brain-computer interface (BCI) performance under different perturbations. *PLOS ONE* **13**, e0191673 (2018).

Jia, K. *et al.* Visual perceptual learning modulates decision network in the human brain: The evidence from psychophysics, modeling, and functional magnetic resonance imaging. *Journal of Vision* **18**, 9 (2018).

Kentridge, R. Spatial attention speeds discrimination without awareness in blindsight. *Neuropsychologia* **42**, 831–835 (2004).

King, J.-R., Pescetelli, N. & Dehaene, S. Brain Mechanisms Underlying the Brief Maintenance of Seen and Unseen Sensory Information. Neuron 92, 1122–1134 (2016).

Kleih, S. C. *et al.* Out of the frying pan into the fire—the P300-based BCI faces real-world challenges. in *Progress in Brain Research* **194**, 27–46 (Elsevier, 2011).

Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., & Broussard, C. (2007). What's new in psychtoolbox-3. *Perception **36***(**14**), 1-16 (2007).

Koch, C., Massimini, M., Boly, M. & Tononi, G. Neural correlates of consciousness: progress and problems. *Nature Reviews Neuroscience* **17**, 307–321 (2016).

Koch, C. *The quest for consciousness: a neurobiological approach*. (Roberts and Co., 2004).

Koch, C. What Is Consciousness? *Nature* **557**, S8–S12 (2018).

Koivisto, M. & Revonsuo, A. Event-related brain potential correlates of visual awareness. *Neuroscience & Biobehavioral Reviews* **34**, 922–934 (2010).

Kouider, S. & Dupoux, E. A functional disconnection between spoken and visual word recognition: evidence from unconscious priming. *Cognition* **82**, B35–B49 (2001).

Kouider, S., Dehaene, S., Jobert, A. & Le Bihan, D. Cerebral Bases of Subliminal and Supraliminal Priming during Reading. *Cerebral Cortex* **17**, 2019–2029 (2007).

Lamme, V. A. F. Why visual attention and awareness are different. *Trends in Cognitive Sciences* **7**, 12–18 (2003).

Lamme, V. A. F. Towards a true neural stance on consciousness. *Trends in Cognitive Sciences* **10**, 494–501 (2006).

Lamme, V. A. F . Separate neural definitions of visual consciousness and visual attention; a case for phenomenal awareness. *Neural Networks* **17,** 861–872 (2004).

Lamme, V. A. F. & Roelfsema, P. R. The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences* **23,** 571–579 (2000).

Lamme, V. A. F. How neuroscience will change our view on consciousness. *Cognitive Neuroscience* **1**, 204–220 (2010).

Lamme, V. The Crack of Dawn. *Open MIND* (2015).

Lau, H. & Rosenthal, D. Empirical support for higher-order theories of conscious awareness. *Trends in Cognitive Sciences* **15**, 365–373 (2011).

Luck, S. J. An Introduction to the Event-Related Potential Technique, Second Edition. (MIT Press, 2014).

Mak, J. N. & Wolpaw, J. R. Clinical Applications of Brain-Computer Interfaces: Current State and Future Prospects. *IEEE Reviews in Biomedical Engineering* **2**, 187–199 (2009).

Marchetti, M. & Priftis, K. Effectiveness of the P3-speller in brain-computer interfaces for amyotrophic lateral sclerosis patients: a systematic review and meta-analysis. *Frontiers in Neuroengineering* **7,** (2014).

Marti, S. & Dehaene, S. Discrete and continuous mechanisms of temporal selection in rapid visual streams. *Nature Communications* **8,** (2017).

Marti, S., King, J.-R. & Dehaene, S. Time-Resolved Decoding of Two Processing Chains during Dual-Task Interference. *Neuron* **88,** 1297–1307 (2015).

Naccache, L. & Dehaene, S. Unconscious semantic priming extends to novel unseen stimuli. *Cognition* **80**, 215–229 (2001).

Naccache, L. Why and how access consciousness can account for phenomenal consciousness. *Philosophical Transactions of the Royal Society B: Biological Sciences* **373**, 20170357 (2018).

Nakamura, K., Dehaene, S., Jobert, A., Le Bihan, D. & Kouider, S. Task-specific change of unconscious neural priming in the cerebral language network. *Proceedings of the National Academy of Sciences* **104**, 19643–19648 (2007).

Owen, A. M. et al. Detecting Awareness in the Vegetative State. Science 313, 1402–1402 (2006).

Owen, A. M. et al. Using Functional Magnetic Resonance Imaging to Detect Covert Awareness in the Vegetative State. Archives of Neurology 64, 1098 (2007).

Pitts, M. A., Martínez, A. & Hillyard, S. A. Visual Processing of Contour Patterns under Conditions of Inattentional Blindness. *Journal of Cognitive Neuroscience* **24**, 287–303 (2012).

Pitts, M. A., Metzler, S. & Hillyard, S. A. Isolating neural correlates of conscious perception from neural correlates of reporting one's perception. *Frontiers in Psychology* **5**, (2014).

Pitts, M. A., Padwal, J., Fennelly, D., Martínez, A. & Hillyard, S. A. Gamma band activity and the P3 reflect post-perceptual processes, not visual awareness. *NeuroImage* **101**, 337–350 (2014).

Polich, J. Updating P300: An integrative theory of P3a and P3b. *Clinical Neurophysiology* **118**, 2128–2148 (2007).

Prinz, J. J. *The conscious brain: how attention engenders experience*. (Oxford University Press, 2012).

Raymond, J. E., Shapiro, K. L. & Arnell, K. M. Temporary suppression of visual processing in an RSVP task: An attentional blink? *Journal of Experimental Psychology: Human Perception and Performance* **18**, 849–860 (1992).

Rensink, R. A., O'Regan, J. K. & Clark, J. J. To See or not to See: The Need for Attention to Perceive Changes in Scenes. *Psychological Science* **8**, 368–373 (1997).

Schelonka, K., Graulty, C., Canseco-Gonzalez, E. & Pitts, M. A. ERP signatures of conscious and unconscious word and letter perception in an inattentional blindness paradigm. *Consciousness and Cognition* **54**, 56–71 (2017).

Schendan, H. E. & Ganis, G. Electrophysiological Potentials Reveal Cortical Mechanisms for Mental Imagery, Mental Simulation, and Grounded (Embodied) Cognition. Frontiers in Psychology 3, (2012).

Scholte, H. S., Witteveen, S. C., Spekreijse, H. & Lamme, V. A. F. The influence of inattention on the neural correlates of scene segmentation. *Brain Research* **1076,**106–115 (2006).

Sehatpour, P., Molholm, S., Javitt, D. C. & Foxe, J. J. Spatiotemporal dynamics of human object recognition processing: An integrated high-density electrical mapping and functional imaging study of "closure" processes. NeuroImage **29**, 605–618 (2006).

Sergent C, Naccache L. Imaging neural signatures of consciousness: 'What,' 'when,' 'where,' and 'how' does it work. Arch. Ital. Biol. **15**, 91-106 (2012).

Sergent, C. *et al.*Cueing Attention after the Stimulus Is Gone Can Retrospectively Trigger Conscious Perception. *Current Biology* **23,** 150–155 (2013).

Sergent, C. The offline stream of conscious representations. *Philosophical Transactions of the Royal Society B: Biological Sciences* **373,** 20170349 (2018).

Sergent, C., Baillet, S. & Dehaene, S. Timing of the brain events underlying access to consciousness during the attentional blink. *Nature Neuroscience* **8,** 1391–1400 (2005).

Sergent, C., Ruff, C. C., Barbot, A., Driver, J. & Rees, G. Top–Down Modulation of Human Early Visual Cortex after Stimulus Offset Supports Successful Postcued Report. *Journal of Cognitive Neuroscience* **23,** 1921–1934 (2011).

Shafto, J. P. & Pitts, M. A. Neural Signatures of Conscious Face Perception in an Inattentional Blindness Paradigm. *Journal of Neuroscience* **35**, 10940–10948 (2015).

Sigman, M. & Dehaene, S. Parsing a Cognitive Task: A Characterization of the Mind's Bottleneck. *PLoS Biology* **3,**e37 (2005).

Squires, N. K., Squires, K. C. & Hillyard, S. A. Two varieties of long-latency positive waves evoked by unpredictable auditory stimuli in man. *Electroencephalography and Clinical Neurophysiology* **38**, 387–401 (1975).

Teufel, C., Dakin, S. C. & Fletcher, P. C. Prior object-knowledge sharpens properties of early visual feature-detectors. *Scientific Reports* **8**, (2018).

Tononi, G., Boly, M., Massimini, M. & Koch, C. Integrated information theory: from consciousness to its physical substrate. *Nature Reviews Neuroscience* **17**, 450–461 (2016).

Tsuchiya, N., Wilke, M., Frässle, S. & Lamme, V. A. F. No-Report Paradigms: Extracting the True Neural Correlates of Consciousness. *Trends in Cognitive Sciences* **19**, 757–770 (2015).

Tsuchiya, N. & Koch, C. Continuous flash suppression reduces negative afterimages. *Nature Neuroscience* **8**, 1096–1101 (2005).

Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences* **36**, 181–204 (2013).

Wilke, M., Mueller, K.-M. & Leopold, D. A. Neural activity in the visual thalamus reflects perceptual suppression. *Proceedings of the National Academy of Sciences* **106**, 9465–9470 (2009).

X. Herman, W. *et al*. A Switch and Wave of Neuronal Activity in the Cerebral Cortex During the First Second of Conscious Perception. *Cerebral Cortex* 1–14 (2017).