Unraveling the Neural Correlates of Consciousness During the Processing of Words

A Thesis

Presented to

The Interdivisional Committee for Biology and Psychology

Reed College

In Partial Fulfillment

of the Requirements for the Degree

Bachelor of Arts

Kavya Basu

December 2016

Approved for the Committee

(Biology and Psychology)

Michael Pitts

Acknowledgments

I would like to thank my brilliant teachers and advisors, Enriqueta Canseco-Gonzalez and Michael Pitts, for their guidance and support. I have learned so much from you both and it's been a pleasure working on this thesis with you. I would also like to thank Corbin, my thesis partner, whose infectious sense of humor kept both our spirits uplifted through all the ups and downs of this year.

Thank you to my friends whose presence and encouragement made this thesis process infinitely more surmountable. Hannah and Andres, you two are the Olympic gold medalists of friends. I am so grateful to have you in my life, thank you for everything. Jen, Shannon, Johnny, Looney, Nick, Mari, and Jenessa—thank you! I am so happy to have known you all at Reed.

Thank you to my family for your unwavering faith in me and deep investment in my happiness. To my sister, Krithi; You're a powerhouse who inspires me constantly even though I'm the older sibling. I love how we turn into cackling fools every time we're together, you're the best. To my parents, thank you for encouraging me to cultivate a sense of wonder and a love for all things mysterious (that's probably why I ended up writing a thesis on consciousness…).

Table of Contents

List of Tables

List of Figures

Abstract

Elucidating the underlying neural mechanisms of consciousness has long been a central goal in neuroscience and psychology. Studies of consciousness have typically relied on reports from participants to link brain activity with their subjective experience. Recently, evidence has been mounting that this dependence on reports results in an overestimation of the neural correlate of consciousness. In this study, a modified version of the inattentional blindness paradigm was utilized to disentangle a few proposed markers of conscious perception without the confounding aspect of immediate report. Word-form stimuli were presented during conditions of unawareness, awareness, and task relevance, while event-related potentials (ERPs) were recorded from brain activity in an attempt to dissociate the neural activity preceding conscious perception from the post-perceptual activity following it. We discovered that the P3b, a well-known candidate for the marker of consciousness, was instead associated with post-perceptual processes and the N200, a marker for visual word-form processing, was indicative of early-stage pre-perceptual processes. The visual awareness negativity was found only during the aware condition and thereby remains a potential neural correlate of consciousness.

Chapter 1: Introduction

The connection between the mind and body has been an intriguing question throughout the history of philosophy and psychology, whether it arose from religious thought in the 6th century BCE or from neuroscience labs today. With the advancements of technology and experimental paradigms to study consciousness, uncovering the neurophysiological processes that accompany conscious perception appears to finally be within our reach. However, several of the paradigms have recently come under question due to the possibility that they overestimate the neural markers of conscious perception. The study detailed in this thesis utilized a methodology that was constructed to avoid the issues of these other paradigms and to isolate the neural correlates of consciousness more accurately. This chapter contains an outline of the history of the field of consciousness, the problems of studying the neural correlates of consciousness, and an introduction to the experiment presented in this thesis.

1.1 The Phenomenon of Consciousness

The nature of our experience of the world has always been a great mystery to be unraveled. Various religions, philosophers, and scientists have attempted to answer the pressing question of our relationship with reality. Ancient Indian philosophies had the concept of *Māyā*, that the world is an illusion—a smokescreen that is not what it appears to be. The phenomenal experiences that we take as an empirical reality are illusory perceptions that we create (Radhakrishnan, 1914). Much later, in the west, Descartes pondered subjects such as thinking, the relationship between our minds and bodies, and our perception of the world. He came up with one of the first theories of mind-body dualism; that the mind is made of a different material than the body. In fact, Descartes pinpointed the locus of this mind substance to the pineal gland in the brain. Spinoza, on the other hand, believed in a more representational theory called mind-body parallelism.

Here he states that the mind is an idea of the body and that the body is represented in the mind (Joregensen, 2014).

In England, Locke followed Descartes in defining consciousness as the awareness of one's own thoughts but not completely as a dualist, as he pointed out that this awareness disappears during varying states of consciousness, such as sleeping or dreaming. Then Leibniz offered one of the first proper theories of consciousness in which he distinguished the perception of the world from apperception, the awareness of perception *i.e.* consciousness. Like Spinoza, he also believed that the mind was representational and that our perceptions were representations of external things in the world (Joregenson, 2014).

The difficulty of defining and studying consciousness is already apparent in these early philosophical investigations of consciousness. It has continued to be a difficult task to this day, with different researchers offering various definitions and taxonomies. Consciousness can refer to vigilance—a state of wakefulness, as opposed to being asleep, anesthetized, etc.—, the pure sensory experience of a piece of information known as phenomenal consciousness, or the experience of information that has entered our awareness and is reportable to others (Block, 1995; Dehaene, 2014). This latter definition of consciousness, known as conscious access, is primarily what is being studied in modern consciousness research since being able to report one's awareness can be used as a measurable variable in scientific experiments (and some would argue that the ability to report on what one is aware of is a fundamental part of being conscious in the first place, even if you cannot report every detail of the experience).

Consciousness was a subject primarily studied in philosophy until the field of experimental psychology developed in the late 19th century. Introspection and the workings of the mind were studied by psychologists such as Wilhelm Wundt, William James, and Edward Titchener (Van Gulick, 2014). However, research into the actual brain was limited at that time. After this period, the rise of behaviorism led to consciousness becoming a taboo subject in psychology for a large part of the 20th century. But as research into cognition grew and the technology to study the brain at a neural level improved, there was a resurgence of interest in the study of consciousness in the 1980s and 90s (Van Gulick, 2014).

David Chalmers (1995) articulated the difference between the “easy and hard problem” in studying consciousness. He referred to the easy problem as mapping out all of the physical processes in the brain that are associated with consciousness. The hard problem is the explanation of *how* exactly these processes become the phenomenal experience we know as awareness of the world. Experimental consciousness research has focused on the easy problem. However, the “easy problem” has proven to not be quite that easy.

Philosophers and scientists who study consciousness today no longer believe in a mysterious substance that comprises the mind or soul, the mind-body (aka mind-brain) problem continues today. The processes of consciousness have been localized to the brain but their exact physiology and even localization within the brain are yet to be uncovered. Still, the work of isolating these neural correlates of consciousness (NCCs) have been greatly abetted by advancements in neurophysiological techniques and the formulation of a number of experimental paradigms.

1.2 Prominent Theories of Consciousness

Much of the research on the neural basis of consciousness has focused on visual awareness. This is due to it being easy to implement in experimental settings and vision being the sensory modality that is most intuitively tied to our experience of the world. The major goal of this research is to uncover the timing of the emergence of visual awareness.

Consciousness theorists are generally split into two camps about this question; sensory theorists and cognitive theorists. Sensory theorists believe that visual awareness is linked to early activation in the visual cortex, which is summarized in recurrent processing theory (Lamme, 2006). Recurrent processing is a four-stage model (Figure 1). It begins with stage 1; shallow feed-forward processing through visual areas. It then proceeds with a deep feed-forward sweep of information about the stimuli from visual to motor areas in Stage 2. Within milliseconds, each area selects perceptual information such as shape, color, position, motion, objects, and faces. In Stage 3, localized recurrent processing allows the exchange of information between and within higher and lower areas, through horizontal and feedback connections. Sensory theorists posit that while information about the stimulus is not accessible for report at this stage, it does achieve phenomenal consciousness. Finally, the widespread recurrent processing to the fronto-parietal network results in reportable conscious experience. Under this framework, all of these stages of processing occur in 100-150 ms and therefore conscious experience of the stimulus should arrive fairly early.

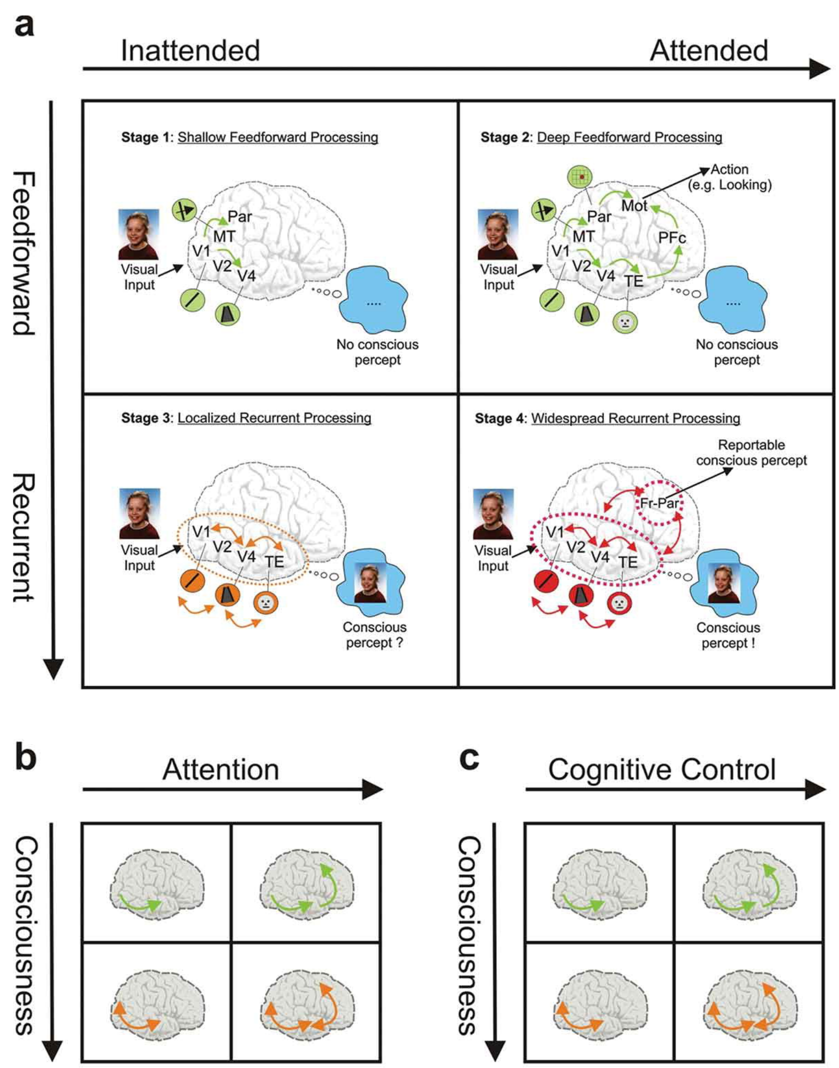

Figure 1. A visual depiction of recurrent processing

a) The four stages of visual processing according to Recurrent Processing theory. b and c) Consciousness is independent from attentional selection as can be seen when feed-forward and recurrent processing is separated and linked to unconscious and conscious processing. Reproduced from Lamme, 2010.

Recurrent processing is in contrast to late theories of consciousness, such as Dehaene's Global Neuronal Workspace Theory (GNWT). This model proposes that conscious processing begins with a feed-forward propagation of stimulus information through a hierarchy of sensory areas (Dehaene, Kerszberg, & Changeux, 1998). If multiple signals converge, it is amplified in a top down manner and becomes maintained by the activity of GNW neurons (Figure 2). This amplification occurs through the "ignition", or a late and sustained firing, of GNW neurons and an increased power of local cortico-thalamic oscillations in the gamma band and long-distance phase synchrony. Crucially, Global Neuronal Workspace theory states that consciousness is all-or-nothing, and that phenomenal consciousness—an in-between stage of consciousness—is not separable from access consciousness (Figure 3). In relation to this view, (and in contrast to Lamme's Recurrent Processing theory) GNWT also does not dissociate attention and consciousness and considers attention to be a prerequisite of consciousness (Dehaene, 2001).

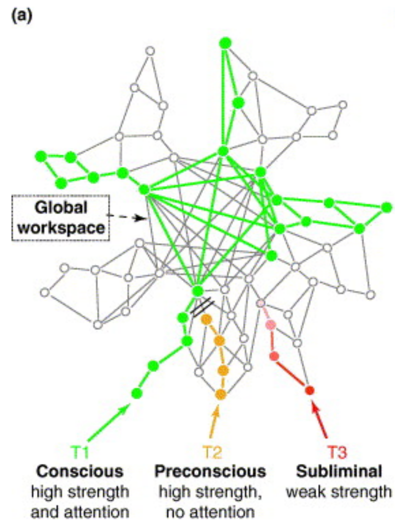

Figure 2. A representation of the Global Neuronal Workspace Model

A visual stimulus (T1) enters consciousness when it activates long-range, sustained, and reciprocal connections a network of GNW neurons. When the stimulus (T2) has good bottom-up strength but poor attention, it does not reach the global workspace and remains preconscious. When the stimulus (T3) has weak activation due poor bottom-up strength it remains subliminal. Reproduced from Dehaene, 2006.

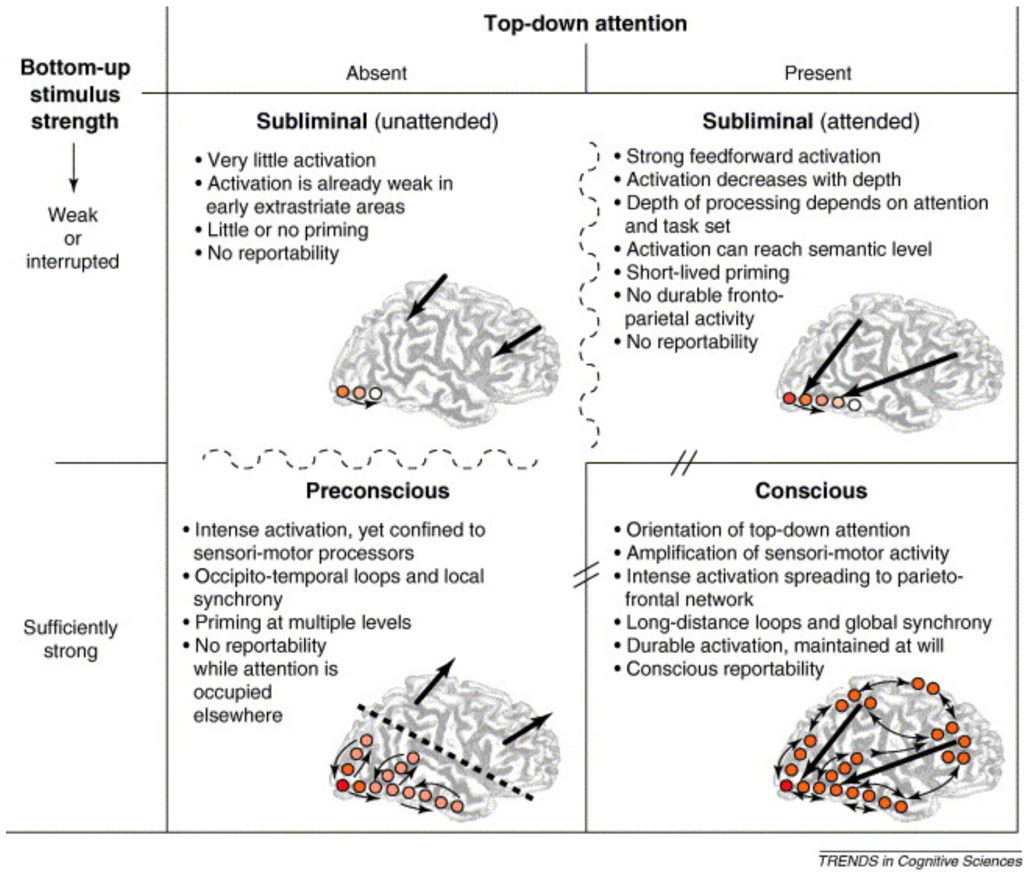

Figure 3. A taxonomy of states of consciousness

States of consciousness as produced by bottom-up stimulus strength and top-down attention. When both attention and stimulus strength are weak, the stimulus is not attended to and subliminal processing with little activation and no reportability occurs. Subliminal processing with slightly stronger activation but no durable fronto-parietal activity occurs when top-down attention is directed to a stimulus with poor bottom-up strength. Preconscious processing occurs when the stimulus strength is high but top-down attention is absent, and there is strong activation and local synchrony which is confined to sensori-motor processors and does not lead to reportability. Conscious processing and reportability occurs when top-down attention is present and the stimulus strength is high, leading to intense activation spreading to fronto-parietal network and reportability. Reproduced from Dehaene (2006).

1.3 Looking for the Neural Correlates of Consciousness

1.3.1 The Three NCC Problem

As mentioned earlier, a proper science of consciousness must be able to explain the relationship between the subjective experience of consciousness and its underlying brain states. The dominant research strategy to identify the neural correlates of consciousness (NCC) has been to connect the behavioral correlates of consciousness to their underlying neural mechanisms.

The neural correlates of consciousness have been defined as "the minimal set of neural processes that are sufficient for a concept percept". It would appear that the most intuitive way to find NCCs would be to contrast brain activity during conditions of awareness and unawareness of identical visual stimuli. Indeed, this method—contrast analysis—has been the dominant approach of studying NCCs. Here, the stimulus conditions are held constant while the consequent variation of conscious perception is measured and the neural correlates are observed (Aru, Bachmann, Singer & Melloni, 2012). This variation in conscious perception can be induced through paradigms that make use of perceptual thresholds, backward masking, change blindness, attentional blink, and ambiguous stimuli (Aru et al., 2012).

For instance, backward masking involves the presentation of a "mask" stimulus immediately after the presentation of the target stimulus, which is briefly presented and not consciously perceived. The stimuli is masked at intermediate latency so that it is seen on half the trials and not seen on half. The idea behind this strategy is that any changes in neural activity would reflect actual markers of consciousness and not merely changes in sensory input; there is always a stimulus followed by a mask. The contrast is between being aware vs. unaware. Correlating changes in perception (from not perceiving to perceiving, or perceiving it as one thing versus another) with changing neural activity would indicate the NCC.

While the contrast analysis method is a logical way to test for NCCs, it has recently come under criticism. Aru et al. (2012) pose the "3 NCC problem". They argue that the contrast of trials with and without conscious perception do not necessarily isolate the NCC proper but may confound it with neural processes preceding and following conscious perception. These processes, the NCC-pr (prerequisites) and NCC-co (consequences), are associated with cognitive components that supplement and are closely linked with consciousness, such as attention, memory, and expectation (Aru et al., 2012). The problem with the backward masking example is that the conditions being compared are one in which participants are unaware of the stimulus and have nothing to report with a condition where they are aware and they can prepare and execute a report. In the first condition, participants are unconscious of the stimuli and in the second, they are aware but have to report it; there is no moment where they are purely aware. This comparison includes not just the neural activity required to generate conscious perception, but also the ability to report it since that is a requirement of the contrast method.

1.3.2 Paradigms for NCC Research

Some of the issues contributing to the 3 NCC problem can be avoided by "no-report" paradigms. These paradigms do not require immediate verbal or button-press reports and instead attempt to infer conscious perception through other means (Tsuchiya, Wilke, Frassle, & Lamme, 2015). These paradigms delay report so that the post-perceptual processes associated with it do not affect measurements of the correlates of awareness. Binocular rivalry, flash suppression, and inattentional blindness are examples of no-report paradigms.

1.3.3 The EEG method

Visual processing occurs quite rapidly, under 150 ms, and while that is not necessarily the timing of visual stimuli entering consciousness, it indicates that the timeframe of interest is very narrow. Therefore, to study the neural correlates of consciousness it is important to use a method with high temporal resolution. The event-related potential (ERP) is a useful measurement due its millisecond resolution. ERPs are "electrical potentials generated by the brain that are related to specific internal or external events" (Luck, 2012). They are measured by electroencephalography (EEG), a method of monitoring the electrical activity of the brain by placing electrodes on the scalp with a conductive gel in between the electrode and skin for the purpose of achieving a stable connection.

Neurons transmit signals to each other through the process of neurotransmission. This involves the release of a neurotransmitter from one neuron's axon and its binding to the receptors on the dendrite of a secondary neuron. The neurotransmitter's release is generated by an action potential, a rapid fluctuation of the cell's electric membrane potential, and its binding to the receptors on the postsynaptic cell results in a postsynaptic potential. These longer-lasting postsynaptic potentials can occur in large numbers at the same time as opposed to action potentials, which occur one at a time and are transient. Postsynaptic potentials can summate and be conducted through the skull and scalp due to the electrical dipole they create, and hence can be recorded as ERPs.

ERPs are obtained from the EEG recording by averaging together several segments of the recording that were time-locked to certain events. These events could be stimuli, responses, decisions, etc. Averaging together several trials that are all time-locked to the same event isolates the brain activity that is associated with the stimulus from the unrelated EEG activity, as well as non-neural sources of electrical noise. Once the averaged ERPs are extracted, the average of one set of trials can be compared to the average of another to determine whether the brain activity differed by condition. This process can be visualized in Figure 4 (Luck, 2012).

Unlike another prominent measurement of neural activity, the BOLD (blood oxygen level dependent) signal in fMRI, ERPs are a direct measurement of neurotransmission mediated brain activity as opposed to the indirect, delayed BOLD signal. Consequently, ERPs have the advantage of being instantaneous, accurate to the millisecond level, and indicative of when neurotransmission is occurring. While fMRIs do have better spatial resolution, due to their temporal resolution ERPs are the appropriate measurement for studying things like visual processing and awareness, which will be coupled with pretty rapid neural activity.

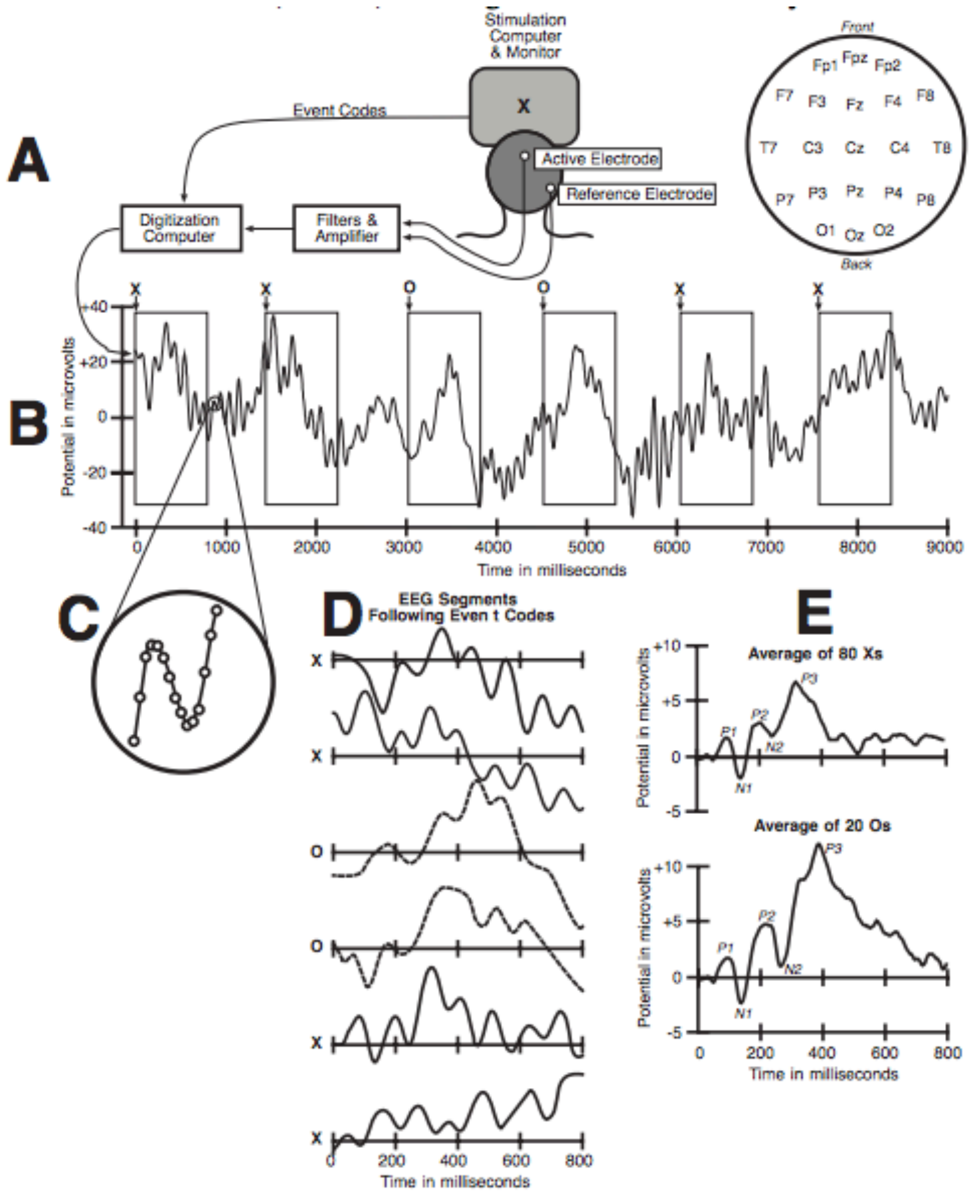

Figure 4. A depiction of the methods used to measure ERPs

A) The locations of electrodes on the scalp with their typical nomenclature. B) A wave-form of the electrical activity recorded by the electrodes with time on the X-axis and voltage on the Y-axis. Event-codes mark the wave-form at certain specified times. C) Points on the wave-form to indicate the sampling rate—evenly-spaced samples of the voltage per second. D) Example EEG segments recorded at each event-code. E) The averaging of the EEG from many trials to produce ERPs. Reproduced from Luck, 2012.

1.3.4 ERPs of interest

Theories of consciousness that propose that the marker of consciousness emerges early, assume that awareness arises at around 200 ms as a consequence of feed-forward and feedback interactions in the visual cortex (Lamme, 2010). Proponents of these theories suppose that conscious perception must arise early from occipitotemporal sites since they believe activity in those areas are the source of visual awareness (Railo, Koivisto, & Revonsuo, 2011). In contrast, late theories believe that visual awareness arises as a consequence of activation of the wide fronto-parietal network and hence the electrophysiological correlate should be observed over a wider area of the cortex in a later time window, around 300-400 ms (Dehaene, Changeux, & Naccache, 2011).

Results from previous ERPs studies show that a larger posterior negativity around 200 ms for stimuli that enter awareness than those that do not (Genetti et al., 2009, Busch et al., 2010, Koivisto et al., 2006). This negativity is termed the “visual awareness negativity” or the VAN. In fact, the VAN and a later positivity after 300 ms—the P3b—usually occur in the same experiments.

GNWT and Recurrent Processing theory make several predictions about the neural correlates of consciousness and highlight the brain areas and ERPs involved. Recurrent Processing theory and other early theories hypothesize that the VAN is correlated with visual awareness and that the P3b is associated with post-perceptual processes (Railo et al., 2011). Proponents of late theories, on the other hand, such as GNWT argue that the VAN could be a preconscious prerequisite of awareness and not awareness itself (Dehaene, 2014). They hypothesize that the P3b is the proper marker of consciousness. The VAN must reflect pre-perceptual processes since some kind of preconscious processing must occur when a stimulus reaches awareness, which does not occur in conditions in which the stimulus is not perceived.

1.4 The Inattentional Blindness Paradigm

A solution to the problem of report corrupting neural measurements of conscious processing is the inattentional blindness paradigm with concurrent EEG recording. The inattentional blindness paradigm relies on the diverting of attention by a distracter task and the absence of participants' knowledge about the target stimulus (Tsuchiya et al., 2015). This is important, because once participants know a stimulus is there they will be able to detect it regardless of dual/distracter tasks. Although unattended stimuli can go through sensory processing, inattentional blindness paradigms block the utilization of attention, cognition, and report, thereby ruling out any post-perceptual processes from determining the marker of consciousness.

However, to study markers of awareness experiments must have a condition of awareness; otherwise, inattentional blindness would only allow us to study unconscious processes. Recent EEG studies have designed a three-phase version of the inattentional blindness paradigm to disentangle pre- and post-perceptual processes associated with visual awareness (Pitts et al., 2012 & 2014). There are three phases: in the first phase, participants perform a task (a distracter task that they think is the primary task) while unexpected and task-irrelevant stimuli appear, which they are hopefully unaware of. After this phase is over, they are then asked (during a delayed report inquiry) if they saw certain stimuli. Participants consequently become aware of the existence of the previously unseen stimuli due to being cued by these questions to their possible presence. In the second phase, they are asked to do the same task again with the only difference being that they are now aware of the stimuli. In the third phase of the experiment, subjects no longer do the distracter task and are asked to provide immediate trial-by-trials reports on the stimuli.

Contrasting phase 2 with phase 1 allows us to isolate NCCs in the absence of immediate report; this can reveal the neural activity associated with awareness of the stimuli without the conflicting influence of processes associated with report. Contrasting phase 3 with phase 2 allows us to isolate neural correlates of report while conscious perception is constant; this can reveal the neural activity associated with post-perceptual, report, and task-related processing.

1.5 The Present Study

1.5.1 Background

Previous research has utilized the modified inattentional blindness paradigm to study NCCs with stimuli made up of simple features. One such study (Pitts et al., 2012) used this paradigm to study the visual processing of contour patterns that would occasionally reconfigure to form simple shapes such as squares and diamonds. As in the modified inattentional blindness paradigm, participants were not aware of the presence of these contour patterns in the first phase of the experiment and were asked to respond to a different task. They were then made aware of the shapes in the second phase while still responding to the same task as in the first phase. In the third phase they were asked to disregard the task and respond directly to the shapes.

This study found an early negativity (220-260 ms), named the ND1, present in all phases of the experiment and a later negativity (~300-340 ms), named the ND2, absent during inattentional blindness but present during both the aware phases. ND1 was hypothesized to be an indicator of an early stage of contour integration while ND2 reflected visual awareness. The authors of this study noted the potential similarity of the ND2 and the VAN reported in previous studies based on its latency, scalp distribution, and polarity.

A similar study (Pitts & Shafto, 2015) investigated the neural markers of conscious face perception. Here, the inattentional blindness paradigm was used to study the awareness of simple line drawings of faces. Subjects performed a demanding distracter task while scrambled lines in the background would sometimes reconfigure to form the eyes, nose, and mouth of a face. They were unaware of this in the first phase, aware in the second, and asked to respond to a facial discrimination task in the third phase.

This study found the N170 (180-220 ms), an early negativity known to be a marker of face-specific processing, and the VAN (260-300) ms—another name for the

ND2 mentioned in the previous study—only in the aware conditions. The P3b (400-600 ms), was found only during the task relevant phase. These results suggest that the N170 is associated with conscious face perception (or attention to faces, a necessary pre-cursor to awareness), the VAN is linked with visual awareness more generally, and the P3b with post-perceptual processes.

This previous research lends evidence to the existence of pre- and post-perceptual processes associated with visual awareness. It also proposes specific ERP components that are representative of different types of neural activity underlying the perception of simple geometric shapes and line drawings of faces. The study detailed here will extend this research by using the same three-phase inattentional blindness paradigm to examine the neural correlates of consciousness for a more complex feature—words.

1.5.2 Visual Word-Form Processing

The Visual Word Form Area (VWFA), located in the left fusiform gyrus of the brain, has revealed sensitivity to specific qualities of visual word-forms (McCandliss et al., 2003). It has been linked with a negative-polarity wave that appears over left posterior visual regions (the location of the VWFA) around 150-200 ms, named the N200. The VWFA has shown selectivity for the visual modality and is not multimodal; spoken words do not elicit neural activity in this region.

There is a limited amount of research on whether the N200 specifically is indicative of the unconscious processing of words. An fMRI and ERP study on the processing of masked words supported the VWFA's ability to unconsciously process visual word-forms (Dehaene et al., 2001). The fMRI results showed activation for the masked words in an area of the left fusiform gyrus that fit with published coordinates of the VWFA. ERPs showed negativity from 212 to 300 ms for masked words and they were found at scalp locations close to the cortical sites where unconscious activation was found with fMRI.

Another masked priming study on automatic word-form processing had similar findings (Grossi & Coch, 2005). Results showed N200s of smaller amplitudes for targets that were preceded by words and nonword primes than for letter strings, false fonts, and neutral primes, signifying that a greater degree of unconscious processing had occurred for the words and nonwords since neural responses are reduced when they have recently been activated.

Despite these results, most of the previous research on the unconscious processing of words has only used paradigms such as masking. Masking has been doubted as a legitimate measurement of unconscious processing since spatial attention is always present and increases the likelihood of the prime entering awareness; in fact, many people are able to report the content of the prime (Kouider & Dehaene, 2007).

However, inattentional blindness is a methodology that diverts attention from the hidden stimulus—the present study is possibly the first to investigate the N200's association with unconsciously presented words. The paradigm used gets around many of the pitfalls of methods like masking.

1.5.3 The Experiment

The present study investigates the neural correlates of the visual awareness of words. The specific questions this study seeks to answer are as follows:

1. Will a marker of visual word form processing emerge as the ND1 did for shapes and the N170 did for faces? Will it appear in all phases including inattentional blindness, or only in the conditions of awareness?
2. Will the visual awareness negativity (VAN) show the same pattern as it did for shapes and faces and arise during aware conditions?
3. Will the P3b show the same pattern as the previous studies, and only appear in the third phase, the aware and task relevant condition?

The results of Question 1 will add to the research on the association between the N200 and visual word-form processing. It will also reveal if the N200 is an indicator of the unconscious processing of word-forms if it emerges when participants are unaware of the words. If Question 2 shows the VAN arising during the second phase, the aware condition, it will fit with its pattern in the NCC studies with shapes and faces and lend further support to the proposition that it is a correlate of consciousness. If Question 3 shows the P3b follows the trend of appearing in the task relevant condition, this will support the theory that it is a post-perceptual process associated with visual awareness thereby contradicting major theories of consciousness that consider the P3b to be a neural correlate of consciousness.

This study utilizes the same three-phase inattentional blindness paradigm as in the previous research, but instead of shapes and faces the hidden stimuli are words. They appear occasionally in a grid of false fonts, which are altered forms of letters of the alphabet so that they are no longer recognizable as letters but serve as a physical control since they are made up of the same segments and intersections.

There was an additional aspect to this experiment: to investigate whether the brain processes semantic information unconsciously. The N400 is an ERP component that is associated with semantic processing and is defined by its larger amplitude for words unrelated versus related to the current semantic context. Previous studies employing an attentional blink paradigm have shown that an N400 is still present during the attentional blink, indicating semantic access despite unawareness (Luck et al., 1996). In this study, it was tested whether the N400 would still be present during inattentional blindness.

Chapter 2: Methods

2.1. Participants

16 healthy adults (13 female, mean age = 21.8 years) with no history of neurological trauma participated in the study. 2 participants were excluded; one for a high number of blinks and being blind during the second phase, and one for noticing the words in the first phase of the experiment. Participants were given five psychology departmental lottery tickets for a chance to win $50. All participants were recruited as volunteers and gave informed written consent prior to participation. Experimental procedures were approved by the Reed College Institutional Review Board.

2.2. Stimuli

Stimuli were presented using Presentation Software (Neurobehavioral Systems Inc., Albany, CA). Participants were seated 70 cm away from a computer monitor on which they viewed the stimuli. The stimuli included scene images, false fonts, related and unrelated words, and color-bisected circle arrays (Figure 5). Stimuli were displayed in the following sequence: first, a scene image was shown for 10 seconds, then a 10x8 grid of false-fonts centered between the four circle array would be displayed for 300 ms, after which the scene would be shown again for 300 ms as well, and then the false-font grid and circle array once again, and so on until the end of the block when the scene image would be shown again for 10 seconds. Displays of the stimuli were separated by an interstimulus interval (ISI) of 700-900 ms (jittered randomly to prevent ERPs being affected due to the expectation of stimuli changes at regular intervals). An example of the stimuli and experiment sequence can be seen in Figure 6.

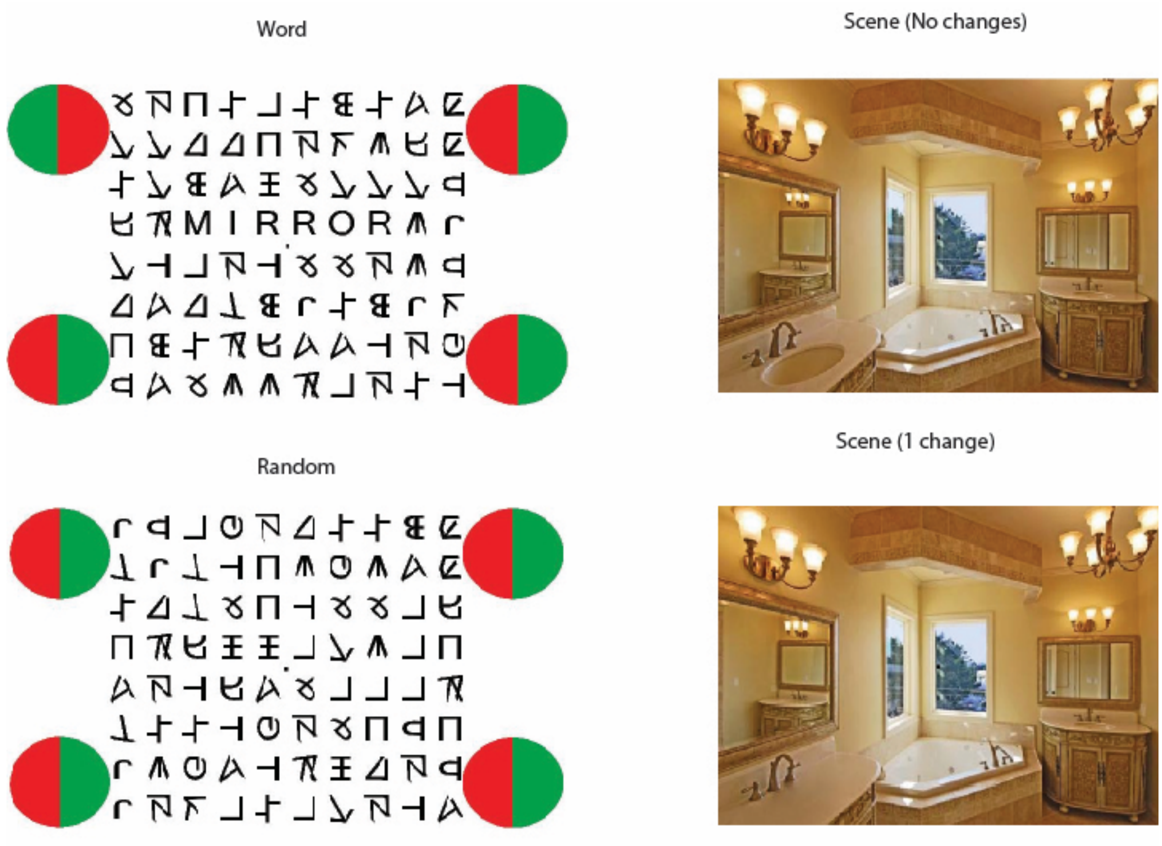

Figure 5. Examples of stimuli

Examples of word/false font grid with circle arrays and the bathroom scene image with and without changes.

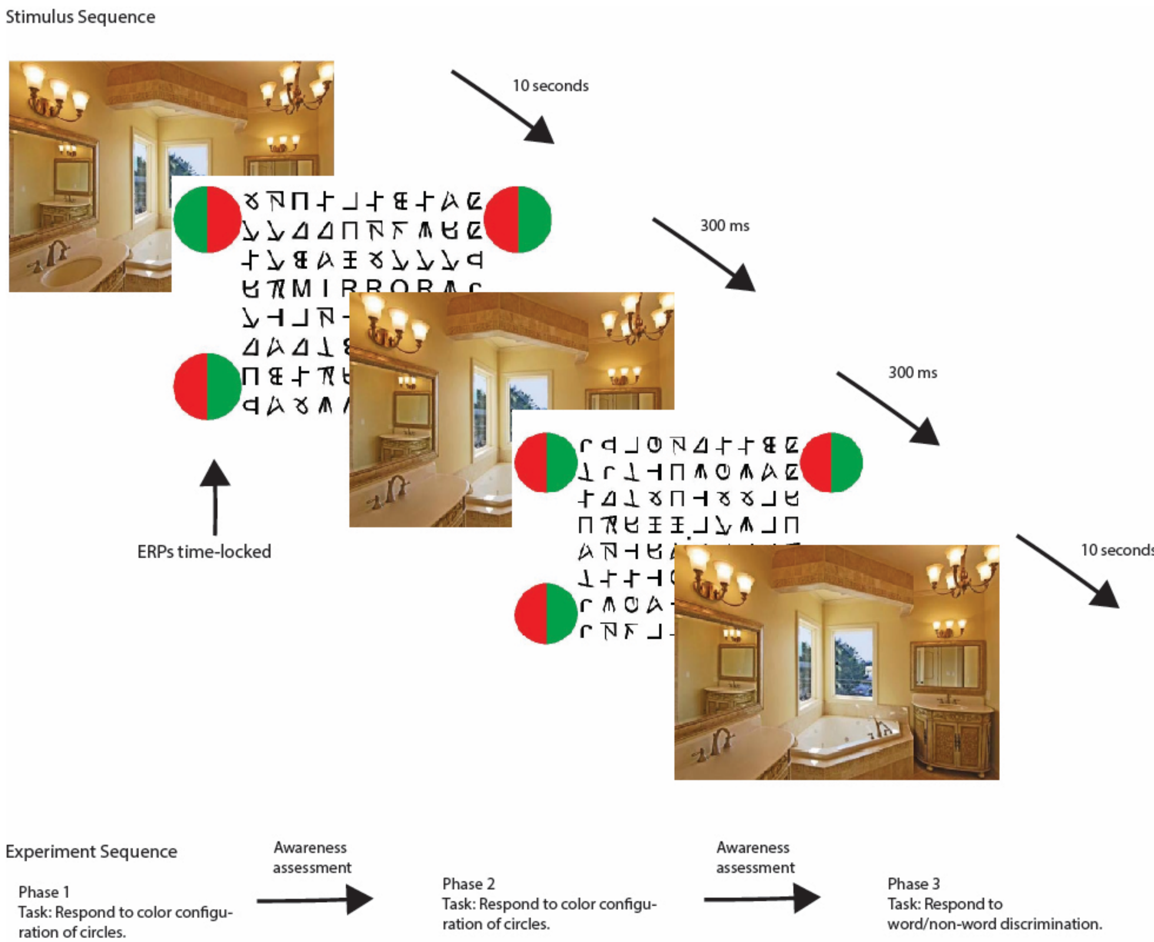

Figure 6. Visual representations of the stimulus and experiment sequences

The false fonts were made in Adobe Illustrator from the 23 letters (3 letters, I, O, and Q, cannot have false fonts) in the Roman alphabet using the Helvetica font. This was done by separating letters into their segments and putting them back together as a false font. The segment lengths and number of intersections were identical to the original letter, but reconfigured unrecognizably. For example, the false font of the letter A would have three segments and three intersections.

There were 11 pictures of scenes, selected from the SUN Database based on their image resolution and ability to adequately represent a certain context (a picture of an airport should be recognizable as an airport). The scenes used were beach, bathroom, airport, bar, art museum, bedroom, hair salon, coffee shop, ski resort, hospital, and classroom. These images were then standardized using Adobe Photoshop. Copies of these scenes were also doctored in Photoshop, to remove some of the items or change their color, for the purposes of the change detection cover task.

For the word stimuli, 10 words related to the category of each context image were selected based on their semantic relatedness to the scene and their frequency of usage in American English. Semantic relatedness was determined using the WUP value (avg = 0.58) from WordNet Similarity for Java. WUP is an indicator of the semantic distance between two concepts and can be a value between 0 and 1. Frequency values (avg = 1523) were taken from the SUBTLEXus word frequency database of 73,000 American English words. 10 unrelated words for each scene category were taken from the 10 other categories, one word from each category. The degree of unrelatedness was again determined by their WUP value (avg = 0.31). Words were between four to eight letters long and made in Illustrator using the same font as the false fonts.

The circle arrays were also made in Adobe Illustrator. There were four bisected circles of the same size, half green and half red. The circles were 100 pixels wide in diameter. The original configuration of the circles was green on the right and red on the left. The configuration with the "switched" circle would have one of the four circles be red on the right and green on the left. They were placed about 3 inches apart in a square formation, with a fixation point in the center.

2.3. Procedure

Participants were told that the study tested visual memory and were asked to detect and remember changes in pictures of scenes that would be displayed intermittently throughout the block. They were told that there could be 0, 1, 2, or 3 alterations to the scene to which they would respond on the number pad on a keyboard. To make the task more difficult they would also be performing a task on a four circle array that was color-bisected, that would be flashed between the images of the scenes. They would have to respond using either of two buttons whether all four circles had the same color configuration or if one of them was different. After receiving instructions, participants were then given three practice blocks to ensure their understanding of the tasks.

Participants with below 60% accuracy on the circle task were given additional practice blocks.

Each block would begin with one of the context images being shown for ten seconds, during which the participant was asked to look around the scene and try and remember its details, after which they were to keep their eyes on the fixation dot in the center of the screen for the rest of the block. After the scene was shown for ten seconds there would be a jittered 700-900 ms ISI and then the circle and false font grid would be displayed for 300 ms followed by another ISI, which would conclude the trial.

On 50% of the trials, a word would be presented either above or below the fixation dot (the remaining 50% were false fonts). Their position on the screen was balanced half and half to the left or right to account for differences in lengths of words. The circles had a 50% chance of being all the same or one different and balanced so that the same number of false fonts and words appeared with a particular alignment of the circle colors. The false fonts were randomly assigned to the grid but balanced so that each false font appeared the same number of times.

There was a total of 40 trials per block, with 20 false fonts and 20 words. In each phase, there were 11 blocks total (one scene per block). Therefore, each phase contained 440 trials, 220 with false fonts and 220 with words. At the end of each block, participants responded to a question asking if they had detected 0, 1, 2, or 3 changes in the scene image throughout the block. There were never actually 3 changes to the photos to prevent participants' attention waning if the maximum number of changes was found.

After the first phase of the experiment, participants were given a questionnaire to fill out, asking whether they had noticed something besides the images, circle arrays, and false fonts appear on the screen. The questionnaire also asked participants, while example images were displayed on the monitor, to rate their confidence in seeing strings of symbols in the grid, the frequency at which they saw each string of symbols, and the point of the experiment at which the symbols started appearing. The images showed the circles and the false font grid with either a string of consonants, a string of numbers, a string of colored squares, an actual word used in the experiment, and just the false fonts.

They then began the second phase which was the same as the first phase, with a different random order of scenes. After the second phase, they were asked to fill out an identical questionnaire the one after the first phase. In the third phase, they were asked to directly respond to the words and perform a discrimination task based on whether there was a word or no word in the center of the grid. EEG was recorded throughout the experiment.

2.4 Awareness Questionnaire

Subjects were divided into two groups, inattentionally blind or noticers, based on their responses to awareness questionnaires given to them after the first phase of the study. 15 of the 16 participants (94%) were inattentionally blind. This was determined by an open-ended question asking if participants saw anything besides the scene images, circles, and random symbols on the screen *i.e.* words, and by three following questions where they were shown five example images of the circle arrays with the false-font grid in the middle and five symbol strings in the center of the grid. The examples were as follows: Example 1 was false fonts as they were in the grid, Example 2 was a string of consonants, Example 3 was colored squares, Example 4 was an actual word from the experiment, and Example 5 was a string of numbers.

Participants were asked to rate between 1 (very confident they did not see the example) and 5 (very confident they did see the example). Participants who scored a 3 or less on the examples with the words *and* who responded that they did not see anything on the open-ended question were considered inattentionally blind. Participants who answered “yes” on the open-ended question or a 4 and above on the word example were deemed not blind. This questionnaire was given again at the end of Phase 2 to ascertain that participants did become aware of the words during the second phase. See Appendix B for an example of the questionnaire.

2.5 EEG Recording

Brain electrical activity was recorded non-invasively from the scalp using customized electrode caps with 64 electrode placements (EASYCAP GmbH, Herrsching, Germany). Electrode impedances were kept below 5 kΩ. Electrode signals were amplified by two modular amplifiers (Brain Vision LLC, Morrisville, NC) and digitized at a rate of 500 Hz. Eye movements were monitored by vertical and horizontal EOG recordings. Each recording session lasted for 2.5-3 hours, which included setup, capping, practice trials, the main experiment, and clean up. ERPs were time-locked to the onset of the critical stimulus presentation (circle/false font grid). Trials with artifacts such as eye blinks, eye movement, or other muscle movements were discarded.

2.6 Data Analysis

EEG data were processed using BrainVision Analyzer software (Brain Products, Germany). EEG was recorded using FCz as a reference, and then re-referenced off-line to the average of the mastoid electrodes. Trials were discarded semi-automatically from analysis if they contained an eye blink (VEOG > 60μV) or eye movement artifact (HEOG > 50 μV), or if any electrodes exceeded predefined signal amplitudes. Participants with fewer than one third (33%) of trials in any given condition after artifact rejection were excluded from analyses to ensure reasonable signal/noise ratio in the averaged ERP waveforms (n = 3). A total of 14 participants were included in the final dataset. ERPs were time-locked to target stimulus onset and baseline corrected at -100 to 0 ms, and low-pass filtered at 30 Hz.

The electrodes and time windows used for analysis were selected based on previously theorized hypotheses. The time window used for analysis of the N200 was 180-240 ms and the electrode was PO8. For the VAN, O10 from 280-340 ms was analyzed. For the P3b, PZ from 300-450 ms was analyzed. Difference waves were calculated by subtract the false font amplitudes from the word amplitudes. The difference wave amplitudes were entered into ANOVAs with the within-subjects factor of phase (1, 2, 3) to determine whether there were mean amplitude differences across phases. Post-hoc T-tests were performed to determine whether these differences were statistically significant. Single sample T-tests against zero were also conducted for each phase to confirm the presence of the ERP components.

Chapter 3: Results

3.1 Behavioral Results

3.1.2 Awareness Questionnaire

Conclusions about whether participants were inattentionally blind or not were made based on their responses to the post-test questionnaire given after the first phase, assessing awareness. If after the first phase, participants indicated that they did not see anything besides the circle arrays, false fonts, and the scene images on the open-ended question and responded that they were less than or equal to 3 (“uncertain”) for seeing the word example, they were considered inattentionally blind.

Fifteen out of our sixteen participants (94%) were inattentionally blind in the first phase. Accordingly, fifteen out of sixteen participants put less than or equal to 3 (“uncertain”) for the word example.

After the second phase, fifteen out of sixteen participants said they saw words on the open-ended question. The one participant who did not see words was excluded from our data analysis (also for their large number of blinks).

3.1.3 Task Performance

a. Accuracy **b. Reaction Time**

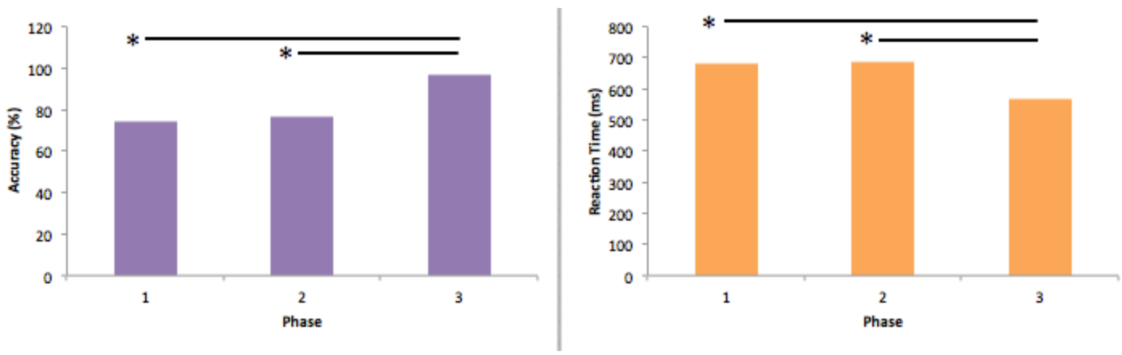

Figure 7. Behavioral results

Mean percent accuracy (a) and reaction time (b) for each phase.

An ANOVA was conducted to check for an effect between phase and reaction times showed a significant effect for phase ($F(1, 13) = 15.66$, $p < 0.00004$) between the first ($M = 679$ ms) and third phase ($M = 564$ ms, $p < 0.013$), and between the second ($M = 687$) and third phase ($p < 0.00007$) indicating that reaction times were faster in the third phase than in the first and second phases.

ANOVA was also conducted with the factors phase and accuracy and also showed a significant effect for phase ($F(1, 13) = 45.74$, $p < 0.000009$). Post-hoc pairwise T-tests revealed a significant difference between the first ($M = 73.9$) and third phase ($M = 97$, $p < 0.00002$) and second ($M = 76.8$) and third phase ($p < 0.00004$). These behavioral results indicate that in general, the word task in the third phase was easier than the circle task in the first and second phases.

T-tests revealed no significant differences between the first and second phase for either reaction time; $t(13) = -0.17$, $p =< 0.8$, or accuracy; $t(13) = -0.66$, $p < 0.5$, signifying that there was no statistically significant effect of difficulty of the circle task and the condition of awareness.

3.2 Electrophysiological Results

3.2.1 Visual Word Form Processing (N200)

Mean amplitude differences from electrode PO8 between 180-240 ms were submitted to a one-way ANOVA. This analysis resulted in no significant amplitude differences across phases ($F(2, 26) = 0.09$, $p < 0.9$). Follow up T-tests confirmed that the N200 was present in all three phases. The mean amplitudes of each phase and the results of the T-tests are provided in Table 1. Grand-averaged ERPs are presented in Figure 8. Difference waves and topographic maps are presented in Figure 9.

N200			
Phase	Mean	t(13)	p
1	-1.24	-2.98	0.01*
2	-1.23	-2.06	0.04*
3	-1.40	-2.52	0.02*

Table 1. Mean ERP difference amplitudes (words – false fonts) for analysis of the N200 and single-sample T-test results.

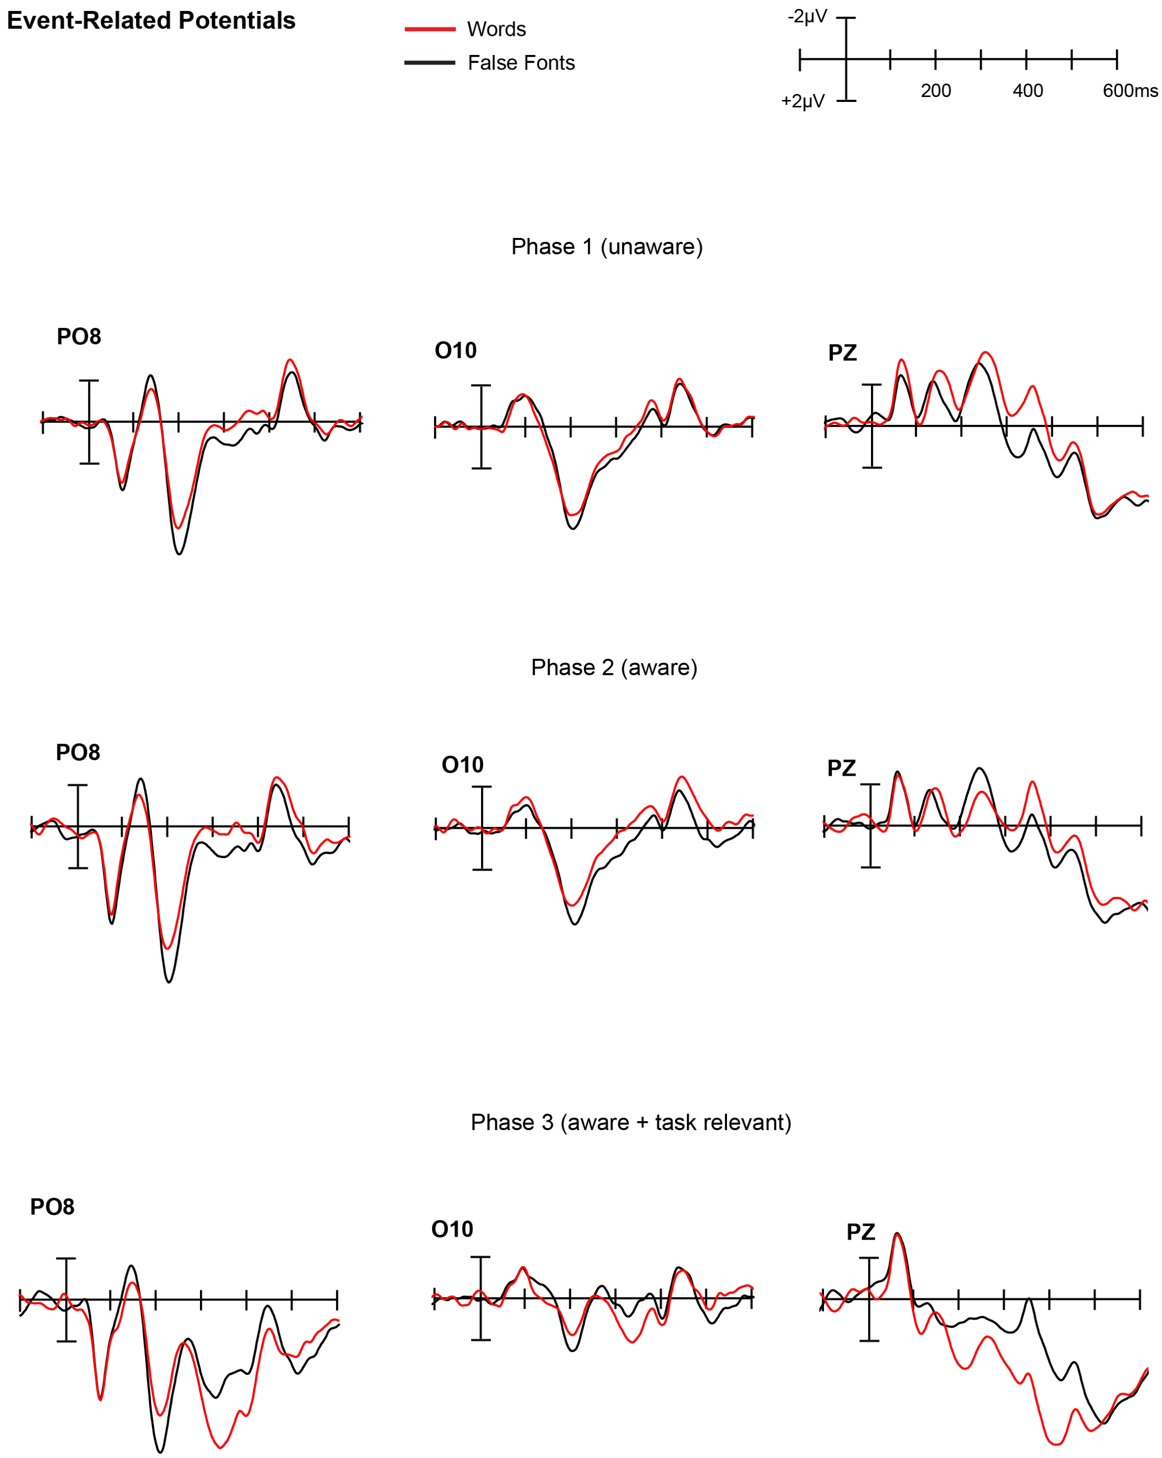

Figure 8. Event-related potentials for word and false font stimuli

Grand averaged ERPs of words and false fonts for each electrode (PO8, O10, and PZ) across the three phases.

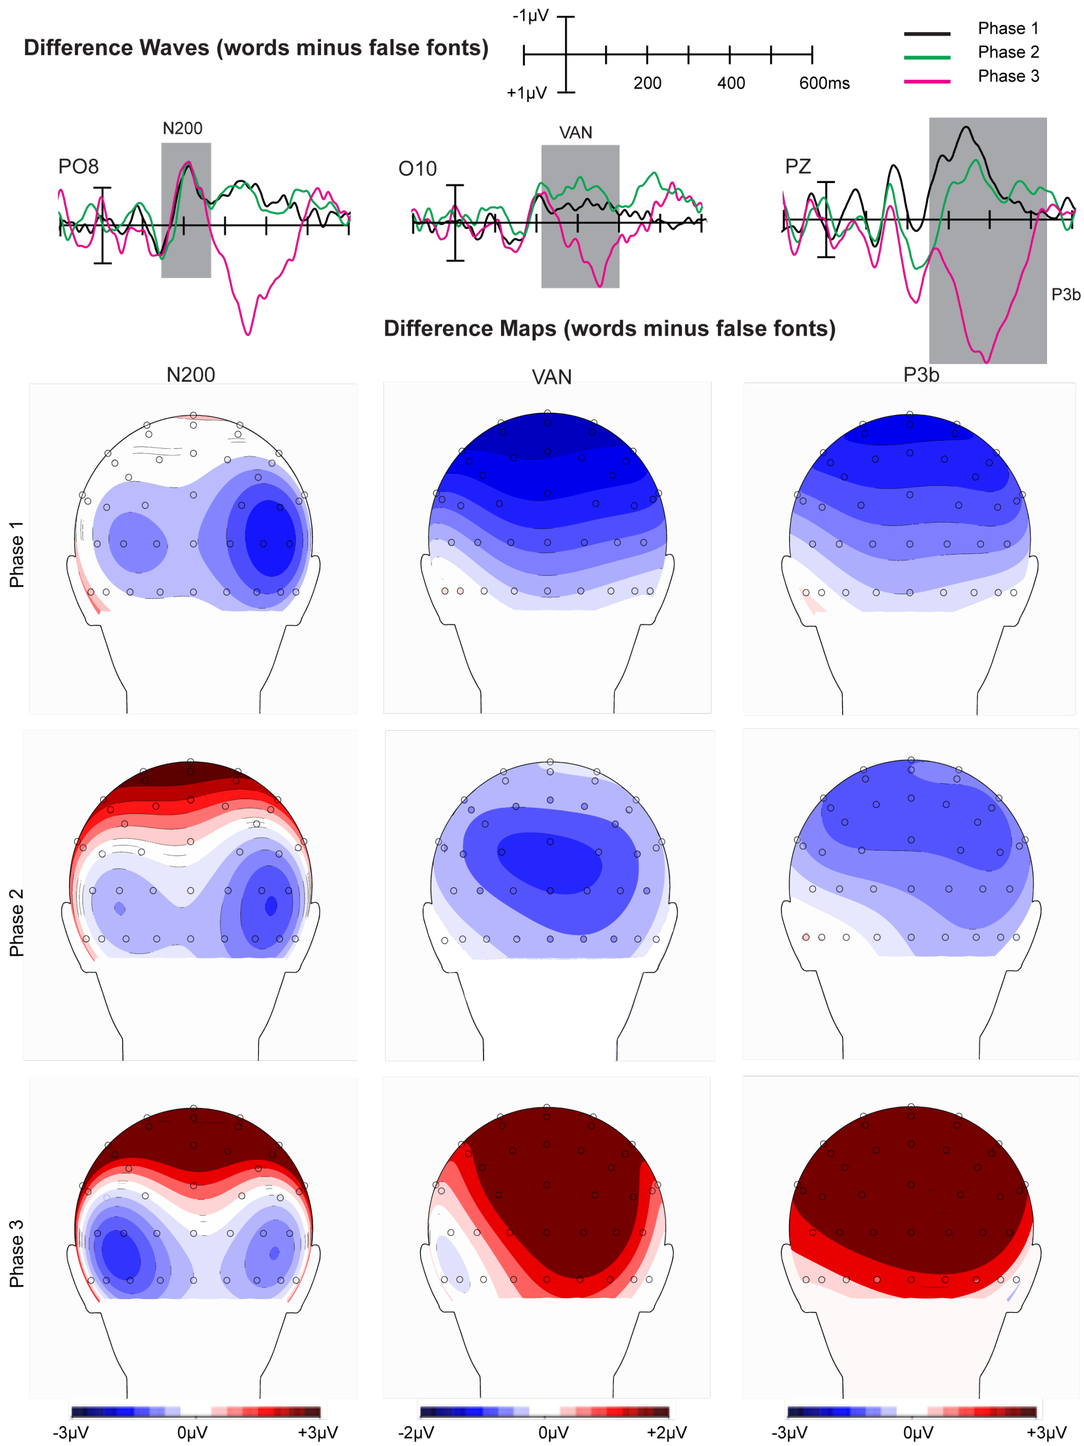

Figure 9. Difference Waves and Maps

At the top are difference waves for each phase and electrode. The N200, VAN, and P3b are highlighted in grey. Below these are difference wave scalp maps for the N200, VAN, and P3b across the three phases.

3.2.2 Visual Awareness Negativity

Mean amplitude differences from electrode O10 between 280-340 ms was submitted to a one-way ANOVA. This analysis resulted in a significant effect of phase ($F(2, 26) = 7.59$, $p < 0.002$). Post-hoc pairwise T-tests revealed a significant difference in amplitudes ($p < 0.0062$) between *aware* and *task relevant* but not between *unaware* and *task relevant* ($p < 0.14$) or between *unaware* and *aware* ($p < 0.45$). Further T-tests revealed that the amplitude of the VAN was significant in *aware* but not in *unaware* or *task relevant*. The mean amplitudes of each phase and the results of the T-tests are provided in Table 2. Grand-averaged ERPs can be found in Figure 8. Difference waves and topographic maps are presented in Figure 9.

VAN			
Phase	Mean	t(13)	p
1	-0.42	-1.64	0.12
2	-0.95	-4.28	0.0008*
3	0.85	2.03	0.06

Table 2. Mean ERP difference amplitudes (words – false fonts) for analysis of the VAN and single-sample T-test results.

3.2.3 P3b

The P3b was analyzed using electrode PZ from the timeframe 300-450 ms. A one-way ANOVA revealed a significant effect for phase ($F(2, 26) = 32.07$, $p < 0.00000009$). Post-hoc pairwise T-tests revealed a significant difference in amplitudes for *unaware* and *task relevant* ($p < 0.00002$), *aware* and *task relevant* ($p < 0.0004$), but not between *unaware* and *aware* ($F(2, 26) = 34.67$, $p < 0.74$). Further single-sample T-tests showed that the mean amplitudes were significant in all three phases; however only the third phase had a positive amplitude indicating that the P3b was present in that phase. The mean amplitudes of each phase and the results of the T-tests are provided in Table 3.

Grand-averaged ERPs are presented in Figure 8. Difference waves and topographic maps are presented in Figure 9.

P3B			
Phase	Mean	t(13)	p
1	-1.53	-5.92	0.00005*
2	-1.03	-2.60	0.02*
3	3.05	5.36	0.00012*

Table 3. Mean ERP difference amplitudes (words – false fonts) for analysis of the P3b and single-sample T-test results.

Chapter 4: Discussion

4.1 Summary of Findings

This study sought to isolate the neural correlate of visual awareness from post-perceptual cognitive processes and to extend the study of NCCs to words. Electrophysiological findings point to a late positivity around 300-450 ms post-stimulus onset as being a correlate of post-perceptual processes associated with consciousness. Although not statistically significant, a negative peak was observed occurring at around 200 ms post-stimulus onset, suggesting it as a marker of visual awareness. A component associated with visual word form processing was present in all three phases.

4.2 General Discussion

4.2.1 N200

Here, we find that the N200 emerged during all phases of the experiment including the unaware condition when comparing the neural activity for words to false fonts. Firstly, this result lends further evidence to the literature associating the N200 with visual word-form processing. Secondly, it indicates that a neural system specified for word-forms is processing words even when they are presented unconsciously. This adds support to the theory that the VWFA is capable of selecting words during conditions of unawareness.

Furthermore, the previous studies investigating the neural correlates of consciousness using shapes found an early negativity (the "ND1") for shapes in all phases of the experiment. The authors proposed that the ND1 reflected an early stage of contour integration. Our findings in relation to the N200 in this study were very similar. Taking into account research that has shown that the N200 is connected to lexical selection rather than semantic processing (van der Brink et al., 2001 & Du et al., 2014), it appears that the N200 is indeed associated primarily with early-stage visual word selection.

4.2.2 VAN

Results from this study show that the VAN was present in the aware phase but not in unaware or aware and task-relevant. Previous studies have shown that typically, the VAN and P3b appear concurrently during phases of awareness. One explanation for our finding is that the immensity of the P3b obscured the VAN. It is also possible that the VAN occurred at an earlier time point in Phase 3 than was analyzed.

Comparing the difference waves across phases only showed significance in the comparison between Phase 2 and 3 but not between Phase 1 and 2, or Phase 1 and 3. Again, this could be because the P3b's amplitude was so huge that the difference between the first two phases and the Phase 3 amplitude that only the comparison between 1 and 3 reached significance.

4.3.3 P3b

The results of this study find the P3b emerging only during the third phase; the task relevant condition. Previous research using contrast analysis methods has found the P3b correlating with conditions of awareness, which would signify visual awareness as emerging late as proposed by theories such as Dehaene's Global Neuronal Workspace Model. The current results contradict this theory, as the P3b only emerged during the task relevant condition and not during the aware condition, in the second phase. These findings support the proposition that the P3b is associated with post-perceptual processes and is not a marker of consciousness.

4.5 Future Directions

While the dissociation of the VAN from the P3b is a promising finding, it would be too hasty to conclude that the VAN must be the neural correlate of consciousness-proper. An extension of these results would be to investigate the properties of the VAN further. One way to do this would be to use other paradigms that modulate awareness to examine how/whether its activity differs.

It would be interesting to continue to examine the N200, its association with the VWFA, and whether it has any specificity for the lexical categories that it activates for. The only category used in this study was nouns. Would the N200 show different responses for other types of words?

4.6 Conclusion

The present study begins to unravel the neural correlates of visual awareness. It puts forth the N200 as a marker for early-stage processing of visual word-forms, designates the P3b as an indicator of task relevance, and leaves the preceding VAN to be a possible marker of consciousness. These results contradict the predictions of a major theory of consciousness; Global Neuronal Workspace Theory. In opposition to GNWT and other late theories of consciousness, it appears that the P3b is likely to be indicative of post-perceptual processes and not consciousness itself. The earlier VAN was found in the awareness condition of the experiment, supporting early theories that put forth consciousness as consequence of activity in sensory areas.

Appendix A: List of scenes and words used

Bathroom	Bar	Art Museum	Classroom	Hospital	Bedroom	Hair salon
bathtub	drinks	artist	teacher	clinic	sleep	hair
toilet	beer	painting	student	doctor	covers	comb
shower	wine	canvas	pupil	nurse	mattress	stylist
towel	liquor	gallery	chalk	surgery	pillow	clippers
mirror	alcohol	exhibit	school	patient	bedframe	spray
sink	cocktail	statue	notebook	sick	sheets	barber
razor	shots	brush	desk	pharmacy	dresser	mousse
soap	bottle	easel	pencil	scrubs	blanket	hairgel
faucet	pint	frame	eraser	needle	pajamas	scissors
tiles	tavern	drawing	ruler	bandage	drapes	blowdry

Coffee shop	Beach	Ski resort	Airport
coffee	sand	snow	airplane
tea	ocean	skiing	airline
espresso	sunshine	gloves	terminal
latte	umbrella	lift	luggage
cream	sail	resort	baggage
sugar	palmtree	lodge	flight
café	water	mountain	boarding
bean	coast	boots	security
mug	shell	coat	suitcase
mocha	shore	scarf	arrivals

Appendix B: Post-test questionnaire

1. Some participants were randomly assigned to conditions in which there were things on the screen other than the scene images, four circles, and the random symbols. While you were completing the task, did you happen to notice anything else on the screen? If so, please describe below. If not, please indicate that.

For the following items, the experimenter will provide examples on the computer screen.

2. Rate how confident you are that you saw each string of symbols during the experiment.

Please use the following scale:
1 = very confident I did not see it
2 = confident I did not see it
3 = uncertain
4 = confident I saw it
5 = very confident I saw it

3. Rate how frequently you saw each string of symbols during the experiment.

Please use the following scale:
1 = never
2 = rarely/ < 10 times
3 = infrequently/10-50 times
4 = frequently/50-100 times
5 = very frequently/ > 100 times

4. At what point in the experiment did you see each of the following symbols start appearing?

Please use the following scale:

1 = training
2 = a previous phase
3 = first half of this phase
4 = second half of this phase
5 = never

5. Any additional comments about this phase may go here.

Bibliography

Aru, J., Bachmann, T., Singer, W., & Melloni, L. (2012). Distilling the neural correlatess of consciousness. *Neuroscience and Biobehavioral reviews, 36(2)*, 737-746.

Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences, 18*, 227-287.

Busch, N.A., Fründ, I., Hermann, C.S. (2010). Electrophysiological evidence for different types of change detection and change blindness. *Journal of Cognitive Neuroscience*, *22(8)*, 1852-1869.

Chalmers, D.J. (1995). Facing up to the problem of consciousness. *Journal of Consiousness Studies*. *2(3)*, 200-219.

Dehaene, S., Kerszberg, M., & Changeux, J.P. (1998). A neuronal model of a global workspace in effortful cognitive tasks. *Proceedings of the National Academy of Science of the United States of America*. *95(24)*, 14529–14534.

Dehaene, S., Naccache, L., Cohen, L., Bihan, D.L., Magin, J.F., Poline, J.B., Riviere, D. (2001). Cerebral mechanisms of word masking and unconscious repetition priming. *Nature Neuroscience, 4(7),* 752-758.

Dehaene, S., Changeux, J.P., & Naccache, L. (2011). The Global Neuronal Workspace Model of Conscious Access: From Neuronal Architectures to Clinical Applications. In S. Dehaene's & Y. Christen's (Eds.) *Characterizing Consciousness: From Cognition to the Clinic, (*p. 55-84) Springer Berlin Heidelberg.

Dehaene, S. (2014). *Consciousness and the Brain: Deciphering how the brain codes our thoughts*. New York: Penguin Books.

Du, Y., Zhang, Q., & Zhang, J.X. (2014). Does N200 Reflect Semantic Processing? –An ERP Study on Chinese Visual Word Recognition. *PLOS One, 9(3)*.

Genetti, M., Khateb, A., Heinzer, S., Michel, C.M., & Pegna, A.J. (2009). Temporal dynamics of awareness for facial identity revealed with ERP. *Brain and Cognition*, *69(2)*, 296-305.

Grossi, G. & Coch, D. (2005). Automatic word form processing in masked priming: an ERP study. *Psychophysiology, 42(3)*, 343-355.

Jorgensen, L.M. (2014). Seventeenth-Century Theories of Consciousness, Retrieved from http://plato.stanford.edu/entries/consciousness/

Koivisto, M., Revonsuo, A., & Lehtonen, M. (2006). Independence of visual awareness from the scope of attention: An electrophysiological study. *Cerebral Cortex*, *16(3)*, 415-424.

Kouider, S. & Dehaene, S. (2007). Levels of processing during non-conscious perception: a critical review of visual masking. *Philosophical Transactions B, 362(1481)*, 857-875.

Lamme, V.A.F. (2006). Towards a true neural stance on consciousness. *Trends in Cognitive Sciences*. *10(11)*, 494-501.

Lamme, V.A.F. (2010). How neuroscience will change our view on consciousness. *Cognitive Neuroscience*. *1(3)*, 204-220.

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

Luck, S.J. (2012). Event-related potentials. In H. Cooper, P. M. Camic, D. L. Long, A. T. Panter, D. Rindskopf & K. J. Sher (Eds.), APA Handbook of Research Methods in Psychology: Volume 1, Foundations, Planning, Measures, and Psychometrics. Washington, DC: American Psychological Association.

Mack, A., Erol, M., Clarke, J., & Bert, J. (2015). No Iconic Memory Without Attention. *Consciousness and Cognition, 40,* 1-8.

McCandliss, B.D., Cohen, L., Dehaene, S. (2003). The visual word form area: expertise for reading in the fusiform gyrus. *Trends in cognitive science, 7(7)*, 293-299.

Pitts, M.A., Martinez, A., & Hillyard, S.A. (2012). Visual processing of contour patterns under conditions of inattentional blindness. *Journal of cognitive neuroscience, 24(2)*, 287-303.

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

Radhakrishnan, S. (1914). The Vedanta Philosophy and the Doctrine of Maya. *International Journal of Ethics. 24(4)*, 431-451.

Railo, H., Koivisto, M., Revonsuo, A. (2011). Tracking the processes behind conscious perception: A review of event-related potential correlates of visual consciousness. *Consciousness and Cognition, 20(3)*, 972-983.

Shafto, J., & Pitts, M.A. (2015). Neural Signatures of Conscious Face Perception in an Inattentional Blindness Paradigm. *The Journal of Neuroscience. 35(31)*, 10940-10948.

Tsuchiya, N., Wilke, M., Frässle, S., & Lamme, V.A.F. (2015) No-Report Paradigms: Extracting the True Neural Correlates of Consciousness. *Trends in Cognitive Sciences, 19(12),* 757-770.

Van den Brink, D., Brown, C.M., & Hagoort, P. (2001). Electrophysiological evidence for early contextual influences during spoken-word recognition: N200 versus N400 effects. *Journal of Cognitive Neuroscience, 13(7)*, 967-985.

Van Gulick, R. (2014). Consciousness. Retrieved from http://plato.stanford.edu/entries/consciousness-17th/#3