Isolating Potential Neural Correlates of Consciousness by Analyzing Pattern Perception Thresholds

---

A Thesis

Presented to

The Established Interdisciplinary Committee of Neuroscience

Reed College

In Partial Fulfillment

of the Requirements for the Degree

Bachelor of Arts

Stella Wroblewski

May 2021

Approved for the Committee

(Neuroscience)

Michael Pitts

Acknowledgements

Thank you to Mom and Jeff for making it possible for me to attend Reed, for still believing in me after my teenage years, and for all that you have done to encourage me to grow and change for the better. to Dad and Tania for inspiring me to seek out beauty and true happiness in life. to my beautiful, talented, and effortlessly wise sister Kayla for being blindly supportive of everything I do and offering love and advice at any time of the day or night. to Katherine, Sonia, Pabu, and Lisa for loving me through all of the good, bad and ugly you have witnessed under our cursed rooves. to Daniel for always giving me a solid connection to my 'roots' in a faraway place and for commiserating past the point of productivity whenever needed. to Jane for confirming my sanity and late night talks. to Kaia for always providing an adventurous and innocent spirit to dull occasions and for being a wonderful listener. to Liam for correcting my endless nonsense, for standing by me during difficult times, and for being handsome. to Paul Currie and Suzy Renn, for showing me real science and being the sole reason I am able to finish Reed in a timely fashion. to Michael and Kevin for all of your help along the way.

I am so thankful for the support and could never have gotten to this point without you.

and to the computer science department, for showing me what I am definitely not missing.

List of Abbreviations

ECoG	Electrocorticography
EEG	Electroencephalography
fMRI	Functional Magnetic Resonance Imaging
GNWT	Global Neuronal Workspace Theory
MEG	Magnetoencephalography
MUA	Mass Univariate Analysis
RPT	Recurrent Processing Theory
NCC	Neural Correlate of Consciousness

Table of Contents

List of Figures

Abstract

The Visual Awareness Negativity (VAN) is an event-related potential (ERP) commonly appearing in occipital and posterior temporal electrode sites on the scalp that has been shown to reliably elicit a significant amplitude difference in aware compared to unaware conditions in the early time range. The present study investigates the neural differences in aware and unaware conditions in regards to identical stimuli as well as fully visible and scrambled images.

To investigate these differences, participants were shown fully visible faces and buildings, individually calibrated threshold faces and buildings, and scrambled images. Following every trial, participants reported whether or not a building or face was perceived. The resulting ERPs were analyzed by comparing fully visible to scrambles and seen threshold to unseen threshold, divided categorically. The VAN, as well as the late component P3b, was shown to significantly differ in amplitude between fully visible faces/buildings and scrambles as well as between seen threshold faces and unseen threshold faces.

Future directions are discussed to further understand the VAN and its categorical differences, as well as efforts towards mitigating the P3b difference through manipulations of probabilities and standardizing spatial presentations of the stimuli.

Introduction

Consciousness

The concept of consciousness is relatively abstract and difficult to strictly define. The current consensus definitions are “subjective experience” or “what it’s like to perceive, feel, or think, from the inside,” or “the brain’s unconscious, embodied enactive, nonconceptual theory about itself” (Cleeremans et al., 2020).

Additionally, the term holds philosophical weight, making it a rich subject for theoretical interpretation and inconsistencies across research. There remains little agreement across fields about how to conceptualize consciousness, therefore, there is much variation in attempts to interpret both theoretical proposals as well as research data. Despite these definitional and conceptual difficulties, Crick and Koch have inspired a new generation of cognitive neuroscientists to set aside these philosophical debates by shifting focus to first identifying neural correlates of consciousness (NCCs), particularly for clearly defined experiences including conscious visual perception.

If we are able to identify and isolate a pattern of brain activity that can be consistently linked with a particular visual percept, such as seeing a face, this brain information can then be differentiated from unconsciously processing an identical face stimulus. This should allow progress to continue and inform the approach to deeper and more complicated “hard problems” of consciousness (Crick and Koch, 1990; Chalmers, 1995).

Technological advancements in ways to measure the brain, such as functional neuroimaging, has engendered new thinking and an opportunity to understand consciousness through empirical neurological data combined with philosophical concepts, rather than from a purely philosophical and theoretical framework. As a result of this shift, and the rapidly expanding depth of consciousness research, several theories of consciousness have gained traction and are increasingly drivers of research to explore consciousness.

Two of the most relevant theories to the current study are the "Global Neuronal Workspace" theory and "Recurrent Processing" theory, as both make detailed predictions regarding the time-course of sensory processing and the particular stage of neural activation most closely linked with conscious perception (Doerig et al., 2020; Förster et al., 2020; Mashour et al., 2020).

Early Versus Late Processing Theories

Global Neuronal Workspace Theory (GNWT) holds that NCCs rely on cortical interactions involving fronto-parietal networks, arising late in time after stimulus onset (post 350ms) and that attention is a necessary component of conscious processing. Recurrent Processing Theory (RPT) states that NCCs result from localized recurrent processing in the sensory cortex, arising earlier in time (pre 150ms) and that cognitive functions such as attention or working memory do not play a role in conscious perception (Pitts et al., 2018).

These two theories can be simplified to the distinction of determining whether consciousness occurs in the time closely following an event or if it occurs more delayed in time, at least for laboratory situations in which time zero of a physical event that is eventually consciously perceived can be precisely measured (Förster et al., 2020).

Manipulations of Awareness

A common approach to studying consciousness is to investigate perceptual awareness, or the conscious experience of sensory stimuli. The two terms, *consciousness* and *awareness,* are commonly used interchangeably in a research setting.

Popular manipulations of awareness can be distinguished by how they manipulate attention. When attention is purposefully directed towards something other than the target stimulus, a stimulus that would otherwise be easily visible may not be perceived. Common paradigms where attention is purposefully directed away include inattentional blindness and attentional blink

(Kim & Blake, 2005). Inattentional blindness manipulates a participant's attention by focusing it on a demanding task, causing the subject to be unaware of an unexpected but otherwise clearly visible stimulus, i.e., the unseen stimulus can be quite salient (high contrast, large, moving, in the center of one's view, presented for a long duration, etc.) while still being undetected due to attention being focused elsewhere (Mack & Rock, 1998).

The attentional blink, on the other hand, involves distracting attention at certain time-points, even when the critical stimulus is expected and relevant to the task (Kim & Blake, 2005). This paradigm involves rapid serial presentation of a sequence of stimuli; when attention is occupied with processing the first target stimulus one's ability to consciously perceive a subsequent target stimulus is drastically diminished if it is presented within several hundred milliseconds of the first (Dehaene & Changeux, 2011).

Awareness of a sensory stimulus can also be disrupted through other means, even if focal attention is fully allocated at the right location and time. Common paradigms include binocular rivalry (Mack & Rock, 1998), masking (Breitmeyer, 1985), and the use of threshold stimuli (Bi & Ennis, 1996).

Binocular rivalry (and related methods such as continuous flash suppression) is a paradigm in which an image presented in one eye can be suppressed from conscious perception by the presentation of a competing image in the other eye, even though it is technically visible (Mack &Rock, 1998).

Visual masking describes a family of paradigms in which otherwise visible stimuli are rendered invisible by presenting additional "masking" stimuli spatially and/or temporally close to the critical stimulus (Breitmeyer, 1985). Threshold detection paradigms involve manipulations of the critical stimuli themselves such that the entire image or a patterned image within a noisy background is degraded to a point at which it is consciously perceived a certain percentage of the time, typically ~50% seen and ~50% not seen (Bi & Ennis, 1996; Dehaene & Changeux, 2011).

Following successful manipulations of awareness via any of the above methods, when combined with concurrent recordings of brain activity, trials where the subject consciously perceives the stimulus can be compared to trials where it was not perceived. Such contrasts may provide insight into the key neural differences between conscious and unconscious states, especially when the physical stimulus, and thus initial stages of sensory processing, are held constant. Upon making these comparisons, researchers have begun to isolate specific neuronal patterns that reliably distinguish between conscious perception and unconscious processing (Dehaene & Changeux, 2011).

Neural Correlates of Consciousness

A neural correlate of consciousness (NCC) is an individual or set of neural signals that is minimally sufficient for indexing a specific conscious experience (Förster et al., 2020). To search for potential NCCs, researchers can compare trials where stimuli are perceived, or not, opening opportunity to isolate conscious from non-conscious processing. The confirmation of a neural signal as a reliable NCC can then be used to evaluate predictions of various proposed theories, for example, whether consciousness arises as a result of early localized processing within sensory regions of the brain or delayed global processing across wide-spread cortical areas (Förster et al., 2020).

There are several different neuroimaging methods applicable to studying consciousness and investigating potential NCCs. The most widely used are functional magnetic resonance imaging (fMRI), magnetoencephalography (MEG), electrocorticography (ECoG), and electroencephalography (EEG). Each of these techniques hold unique advantages and disadvantages in the context of consciousness research. It is important to understand the various applications of each technique to determine which is best suited for an individual experiment and how each can provide varying insight into consciousness research overall.

FMRI

FMRI measures neural activity indirectly by detecting changes in localized cerebral blood flow and oxygenation concentration. When a localized population of neurons was recently active, there are detectable changes in the relative amounts of oxygenated versus deoxygenated hemoglobin. These changes can help researchers localize which brain areas are involved in specific sensory, cognitive, or motor tasks.

The main strength of fMRI is that it is spatially precise (~1 mm), however fMRI provides a much coarser view over time (the blood oxygen level dependent signal peaks at ~5-6 seconds after the onset of a stimulus) compared to the duration of the underlying neural processes (which are on the order of 10-100 msec), and therefore is not practical for use in studies requiring temporal precision (Glover, 2011).

MEG

MEG measures the neuronal activity in the brain by detecting magnetic fields. These magnetic fields are generated from the flow of electrically charged ions through cells via electrochemical properties of neurons. MEG uses sensors to detect and amplify these magnetic fields generated by neurons (University of Washington Institute for Learning and Brain Sciences 2020). MEG has the advantage of high temporal precision, however the spatial precision is not as accurate as fMRI (Puce & Hämäläinen, 2017), and similar to fMRI, the equipment required to acquire MEG data is quite expensive (> 1 million USD).

ECoG

ECoG measures electrical brain signals via electrodes implanted on the surface of the brain, or sometimes deep within the cortex of the brain. ECoG produces very clear output, as it uses electrodes directly on/in the brain, as well as high temporal and spatial resolution. However, it is highly invasive and can only be measured if there is a separate medical reason for neurosurgery, e.g., severe epilepsy (Hill et al., 2012). Accordingly, the placement of electrodes in human brains during neurosurgery is entirely determined by clinical as opposed to research interests, thus the main drawback of ECoG is the limited (and variable across subjects) spatial coverage of measured neural signals.

EEG

EEG records electrical activity in the brain via electrodes placed on the surface of the scalp. Event-related potentials (ERPs) are derived from EEG recordings when a signal-averaging approach is used to associate neural activity with the presentation of a stimulus or general event in time. EEG is noninvasive, has high temporal precision, wide spatial coverage, and is widely available due to its affordability (more than 10 times cheaper than fMRI/MEG). EEG is the best suited imaging technique for the current experiment because the questions being asked require high temporal precision, with spatial coverage being more critical than spatial resolution, and the convenience and non-invasiveness of EEG allows for a larger sample size to ensure generalizability of the results.

With time-resolved methods for recording brain activity, such as EEG, three of the prominent potential NCCs that have been proposed as candidate signatures of visual awareness are the N170, the P3b and the VAN. The N170 is an early occipital-temporal negativity (peaking approximately 160-200 ms following stimulus presentation) that has been shown to correspond specifically to the presence of faces, one of the most important categories of stimuli to humans (Bentin et al., 1996; Schubert et al. 2020). The P3b is a late positive wave with a broad topographical distribution, appearing approximately between 300-600 ms following a stimulus (Polich, 2020). The visual awareness negativity (VAN) is an early/mid latency neural signal occurring about 200-300 ms following stimulus presentation, defined by a relative amplitude difference in aware versus unaware conditions (Förster et al., 2020).

N170

The N170 is an early negativity (in the 140 to 220ms time range) that has been shown to appear when subjects view a stimulus that includes a human face (Harris et al. 2013). In previous literature, the N170 has been shown to be either significantly reduced or totally absent in unaware (or perceptually invisible) conditions (Shafto & Pitts, 2015). The convincing evidence of the N170 as a marker of awareness, sustained in no-report paradigms involving inattentional blindness, leads it to be a high ranking potential NCC (Shafto & Pitts, 2015).

Studies using masking have also confirmed that the N170 is not present in trials with invisible stimuli (Rodríguez et al., 2012). However, an attentional blink study by Harris et al. found this signature to be present in unaware conditions. This study showed the N170 component to be present regardless of reported awareness of the stimulus by the participants. Not only was the N170 present in conditions where the participants were unaware of the stimulus, but there was no significant amplitude difference between aware and unaware trials.

Additionally, this study found that the late phase of processing specific to faces was associated with awareness, where the early stage of specific processing, seen via the N170, was not awareness dependent (Harris et al., 2013). Interestingly, in a recent study by Schubert et al., a special-case participant with a neurodegenerative condition allowed for novel manipulations of awareness. The condition, selective metamorphopsia, disrupts perception of specific letters and digits, as well as clear images embedded within them. The participant elicited a consistent N170 despite any awareness or ability to detect the presence of a face (Schubert et al., 2020).

The literature surrounding this topic is inconsistent in its ability to classify the relationship between the N170 and visual awareness of faces. It remains an open question as to whether the N170 can be deemed as an indicator of conscious perception, highlighting the necessity to further investigate its role.

P3b

Introductory sentence about time range and difference The P3b was long thought to be high on the list of potential NCCs, as it has been shown to robustly appear in aware conditions and disappear, or greatly recede, in unaware conditions (Cohen et al., 2020). However, though the P3b is known to routinely appear in aware states in studies with trial-by-trial reports, several studies have recently noted that the P3b consistently disappears when the stimuli are task-irrelevant, in so-called “no-report paradigms” (Tsuchiya et al., 2015; Cohen et al., 2020). For this reason the P3b is beginning to be widely discounted as a potential neural correlate of consciousness as it cannot reliably isolate aware versus unaware states and is reliant on the presence of post-perceptual processing needed to provide behavioral reports, having little to do with conscious processing per se (Förster et al., 2020).

VAN

Recent research has pointed to the visual awareness negativity (VAN) as "the earliest and most reliable correlate of visual phenomenal consciousness" (Förster et al., 2020). VAN commonly appears in occipital and posterior temporal electrode sites on the scalp. It has been shown to reliably elicit a significant amplitude difference in aware compared to unaware conditions. VAN can begin as early as 100ms after the presentation of a stimulus and end as late as 350ms following (Förster et al., 2020). In an analysis of 30 recent studies looking at the VAN along with other potential NCCs, Förster et al. concluded that 21 found the VAN and only 13 found the LP (late positivity in the range of the P3b) in aware conditions (Förster et al., 2020). Of the three discussed potential NCCs, the consistency of the VAN shows the greatest potential to be a true signature of visual consciousness.

Although an auditory equivalent (auditory awareness negativity) and somatosensory equivalent (somatosensory awareness negativity) have been reported with similar timing but distinct scalp distributions, no previous study has investigated differences in VAN topography for awareness of different categories of stimuli within a given modality (Eklund et al., 2020; Auksztulewicz et al., 2012).

The presence of color areas (V4/V8) (Shapley & Hawken, 2011), face areas (fusiform face area) (Kanwisher et al., 1997), place areas (parahippocampal place area) (Epstein, 1998), and motion areas (V5/MT) (Bayerl & Neumann, 2010) are generally understood, but it is not known whether these areas contribute specifically to the potential NCCs for conscious awareness.

Threshold Perception

Of the paradigms mentioned earlier, there are unique disadvantages to those that purposefully redirect attention, such as attentional blink and inattentional blindness, as they have been found to result in failed or delayed processing due to the displacement of attention. Because they aim to manipulate initial processing, they can have the effect of mitigating early signatures of visual processing.

However, paradigms that require focused attention and often use simple tasks (binocular rivalry, masking, threshold stimuli) are also potentially problematic. These work by reducing bottom-up, stimulus driven information, making the stimulus undetectable to the viewer, regardless of focused attention. But because participants' attention is consistently directed to the stimulus, these paradigms may engender familiarization with the task and stimuli, potentially compromising the ability to successfully manipulate awareness through the course of an entire experiment (Dehaene & Changeux, 2011).

An individual's absolute threshold is defined as the level of a particular stimulus detected in approximately 50% of trials (American Psychological Association, 2020). The main disadvantage of using threshold stimuli is that each subject's perceptual threshold is likely to be different. This is logistically problematic instead of scientifically problematic This disparity means that it is often necessary to employ varied stimuli adjusted to each individual's threshold level. However, this individualization of the stimuli could impact the consistency of the stimulus set across subjects, creating less control vague (Dehaene & Changeux, 2011).

A compelling advantage of presenting threshold stimuli as opposed to mask is that the mask often creates an overlapping neural response in addition to the stimulus itself. When analyzing the data, the mask must be individually subtracted from the masked stimulus presented, so that the data is not impacted by the brain's response to the mask itself. There is an additional possibility that the mask would not translate to a consistent relationship with the blank mask, potentially producing inaccurate conclusions. When attempting to create consciously perceived and unperceived stimuli, it is advantageous to maintain a consistent stimulus set, ideally only varying subjective awareness (Förster et al., 2020). Presenting threshold stimuli allows for this, as it is possible to present the same stimulus that is sometimes visible and sometimes not, without the addition of a mask.

The current study uses threshold stimuli, as this paradigm is most adequately suited to the goals of the experiment. Absolute thresholds are understood to be the smallest level of the stimulus detectable in 50% of trials. Already defined this (maybe delete sentence above) Pattern perception thresholds present a salient stimulus, regardless of whether or not it is perceived, allowing for more robust early visual processing than a weaker stimulus. Threshold stimuli produce cleaner ERPs and more reliable results than masking, and present more robust early visual processing than paradigms that redirect attention. This aids in the investigation of early neural signatures such as VAN and N170, and provides the opportunity to compare unseen trials where a stimulus is present to trials where there is no visible stimulus, yielding insight into the relationship between conscious perception and awareness (Dehaene & Changeux, 2011).

Current Study

A previous study conducted in the Reed College SCALP lab investigated neural signatures of visual awareness using report and no-report paradigms. Participants were shown well-above threshold faces, scenes, or scrambled stimuli in both a report and a no-report condition. In the report condition, subjects were asked to report if they saw a face, scene, or scramble. The results confirmed previous findings detailed above, that the P3b is only sustained in the report condition. Additionally, they found a robust N170 and VAN in trials presenting faces in both the report and no-report condition (more detailed data will be presented following data analysis). However, the “seen” faces (and scenes) were physically different than the “unseen” faces (and scenes) which were just scrambled noise patterns. Thus, the NCC were likely overestimated in this previous study, as some of the neural differences were probably due to early sensory processing differences between the physically distinct classes of stimuli.

To expand on the findings in the report and no-report conditions, we sought to further understand and possibly rule out potential NCCs isolated in the previous study. It was determined that this might be achieved by expanding the experiment to look at threshold stimuli (physically identical in the "seen" and "unseen" trials) to look for evidence that these potential neural markers are linked with conscious visual perception.

The previous study showed a VAN, but due to the nature of the design it was only possible to compare highly visible and never visible stimuli that were physically different. However, the standard method to measure the sVAN is by comparing seen versus unseen trials using the same physical stimulus. Based on the review by Förster et al., the VAN is the most consistent potential early marker of consciousness (Förster et al., 2020). This study will specifically look more closely at this potential NCC and, by utilizing two categories of stimuli – faces and buildings – will allow for an investigation into whether the spatial distribution of the VAN differs when the content of visual consciousness differs.

The stimulus set consisted of 10 different face and building (5 in each category) images embedded in varying levels of noise. The noise levels were manipulated to present 400 images that were exclusively noise (never seen), 400 images that were low noise (always seen), and 800 images that were embedded in the amount of noise tailored to each subject's perceptual threshold (ideally seen in half of trials).

To find each subject's threshold perception, participants engaged in a preliminary portion of the experiment in which they were repeatedly asked to report whether they saw something, or not, for varying noise levels in each individual stimulus. Based on the results in this preliminary test, each participant yielded a unique range approximately representing their individual threshold for each different face and building presented. This range accounts for the level of noise in which the threshold stimuli was embedded for each subject in the experiment.

Throughout each block, the percentage of seen and unseen reports of threshold stimuli were recorded. At the end of each block, if the percentage of seen trials exceeded roughly 60%, the level of noise was lowered accordingly, and if seen trials dropped below 40% visibility, the level of noise was raised accordingly. To mitigate the effect of participants learning the stimuli in the course of a single block, and therefore potentially being able to quickly increase visibility, stimuli were presented behind a virtual window. Behind the window, a consistently positioned frame, the position of the stimuli was spatially jittered across trials. This allowed for variation in the stimuli while maintaining consistency in the field of vision, noise patterns of each image, and the images themselves.

After each trial, subjects were asked to report if they "saw something" or "saw nothing." This was a change from the previous study, removing emphasis from discriminating the stimuli identifying the categorical differences and focusing on perception, or the absence thereof.

In comparing the always visible to stimulus absent (noise-only) trials, we can isolate a broad set of potential NCCs that may represent the difference between seen versus unseen conditions, and directly compare data from this experiment to the previous study mentioned above. Furthermore, uniquely in this experiment, we can compare seen to unseen threshold stimuli to determine which subset of these potential NCCs consistently indicates awareness of the stimuli when the physical input to the visual system is held constant.

If we are able to successfully target participants' thresholds to about 50% seen and unseen, we expect to see the P3b and VAN differentiating aware from unaware conditions and the N170 in aware and possibly unaware trials with faces. By comparing the results of each portion of the experiment, we can gain insight into the theories of consciousness regarding timing as well as increasing understanding of the potential neural correlates of consciousness.

Methods

Participants

A total of 25 Reed College students ages 18-22 with normal or corrected vision and no history of traumatic brain injury participated in this study. All subjects received $20 for approximately two total hours of participation, funded by the Reed College Psychology Department. 10 participants were excluded for reasons relating to behavioral data patterns or problems with the EEG data, leaving a total of 15 participants included in the final analyses below.

Individualizing Threshold

In the first portion of this experiment, participants were shown images of faces and buildings embedded in varying levels of visual noise (figure 2.1). The task was to report whether or not they perceived something in the image presented (either a face or a building). When the subject perceived a face or building within the last image presented, they were instructed to select "yes" and if not to select "no" by pressing corresponding keys on the keyboard. Based on their responses, a staircase algorithm converged to the level of noise that the participant perceived 50% of the time for each individual face or building presented (10 total - 5 faces and 5 buildings). The program then produced a list of transparency values corresponding to each (Shuster, 1978).

Stimuli and Procedure

In the main experiment, participants were shown the same group of stimuli, broken into 400 clearly visible faces and buildings, 400 scrambles, and 800 threshold images tailored to the participants' threshold for that individual stimulus (10 total). Each stimulus was presented for 50ms followed by a variable 1350-1550ms blank in which participants were expected to provide their response for the previous stimulus by pressing keys corresponding to a yes or no response.

The experiment consisted of 20 blocks, each containing 80 trials. Following each block, the experimental program would assess the detection rate of threshold stimuli within the previous block. If the participant was reporting either above 60% or below 40% detection, the program would decrease or increase, respectively, the visibility of all threshold stimuli by a factor of the standard deviation between the original visibilities provided by the first portion of the experiment. Figure 2.1 shows an example of a highly visible face and building at 50% visibility (left), a threshold face and building at 25% visibility (middle), and a never visible scramble at 0% visibility (right).

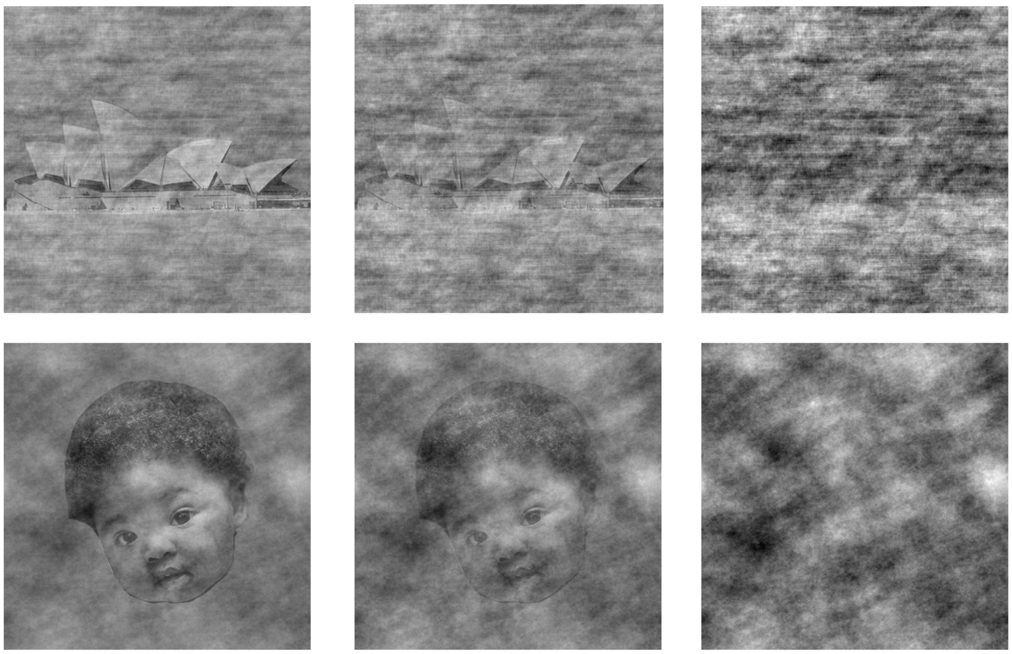

Figure 2.1 Categories of Stimuli
(Top Left) Fully Visible Building. (Left Bottom) Fully Visible Face. (Center Top) Threshold Building. (Center Bottom) Threshold Face. (Right Top) Never Visible Scramble Created From Building. (Right Bottom) Never Visible Scramble Created From Face.

EEG Recording

All participants were fitted using a 64-channel electrode cap, diagrammed by figure 2.2. An electrode placed below the left eye was used to detect blinks (VEOG) and two electrodes were placed next to each eye to detect horizontal eye movements (HEOG). The impedance levels of all electrodes were kept below 5KΩ.

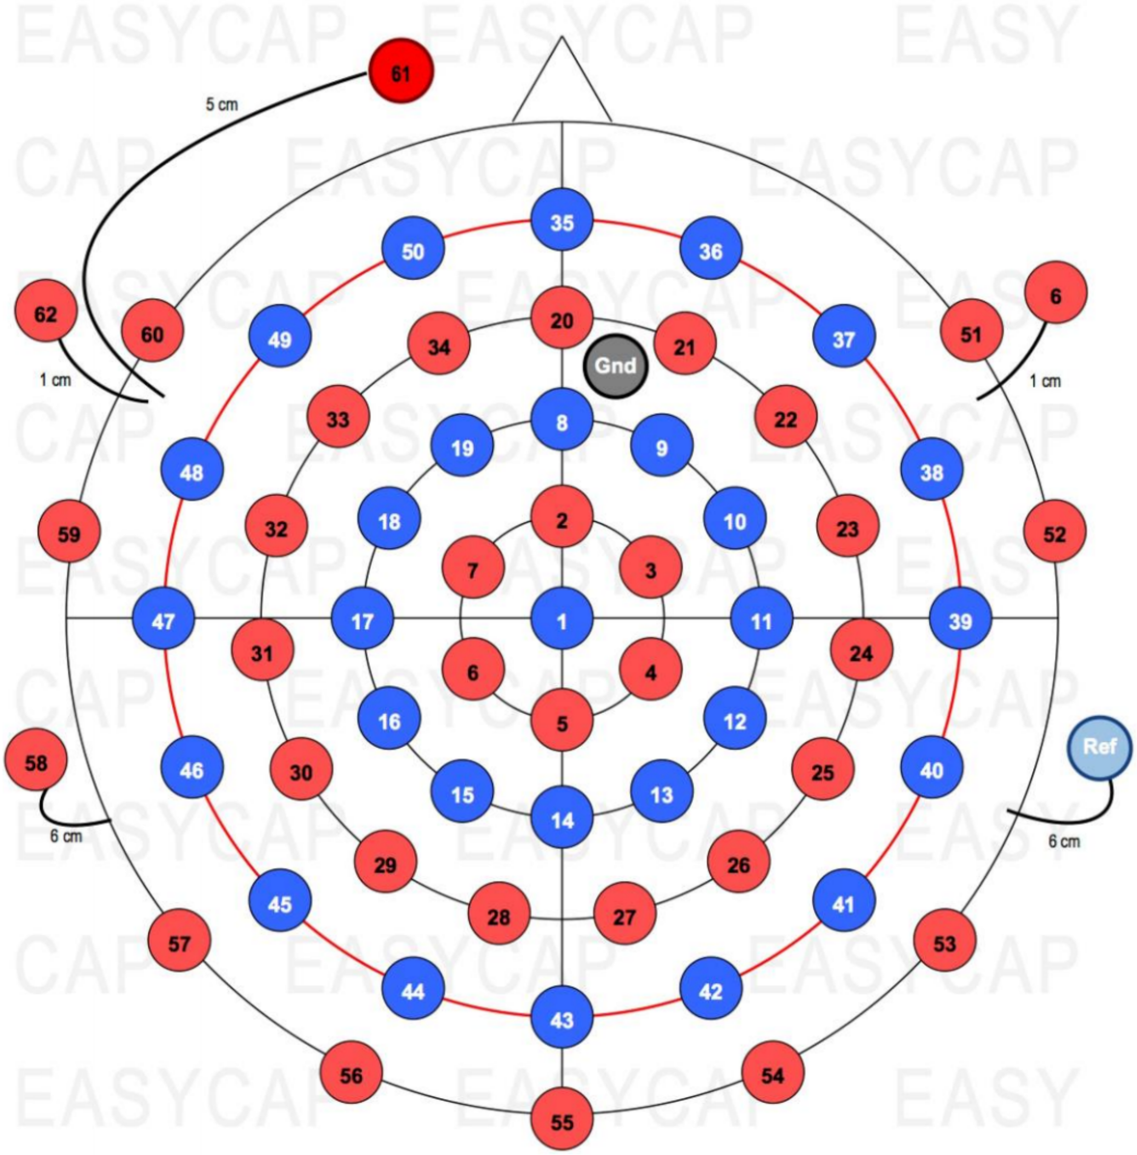

Figure 2.2 64 Channel Electrode Locations
A diagram of the 64 channel EEG cap used. Channel 61 (VEOG) was used to detect blinks and channels 62 and 63 (HEOG) were used to detect eye movements.

EEG Analysis

All EEG data were processed using BrainVision Analyzer software (Brain Products, Germany). Artifacts including blinks, eye movements, muscle related noise, etc. were rejected semi-automatically (on average 24%). EEG was recorded using a right mastoid electrode as a reference, and re-referenced in the analyzer to the average reference. ERPs were time locked to the onset of the stimuli. Participants were excluded if they had less than 100 trials remaining in either the fully visible, threshold faces, threshold buildings, or scramble categories after artifact rejection.

Mass Univariate Analysis

Mass Univariate Analysis Toolbox in MATLAB was used for the statistical analysis of the EEG data. Difference waves obtained by subtracting unseen scramble trials from seen highly visible trials as well as unseen threshold trials from seen threshold trials were subjected to repeated measures, two-tailed t-tests for each condition 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 0 and 500ms for all 59 scalp electrodes (excluding HEOG and VEOG) were included in the analyses.

Results

Exclusion Criteria

Four subjects were excluded due to excessive EEG artifacts, resulting in less than 100 trials per condition in the averaged ERPs of interest. Participants were additionally excluded if their individual false positive (reporting they saw a stimulus when only noise was presented) or false negative (reporting they didn't see a stimulus when a clearly visible stimulus was presented) rates were greater than 30%. Participants were also excluded if their threshold detection rates for either faces or buildings were above 70% or below 30%. Six subjects were excluded based on these behavioral criteria. A total of 10 subjects were excluded based on behavioral or EEG exclusion criteria, leaving 15 subjects included in the final analyses.

Behavioral Results

Within the final group of subjects (N=15), the average rate of detection of all threshold stimuli was 48%, while the detection of threshold faces was 47% and threshold buildings was 49%. In the fully visible condition, the average rate of detection for all stimuli was 97%, while the detection of faces was 98% and buildings was 96%. In the never visible scramble condition, the average rate of detection (false positives) was 8%, meaning on average, subjects reported not seeing a face or building 92% of the time when the stimulus contained only scrambled "noise." Average detection rates across conditions can be seen in figure 3.1. Detection rates for each individual subject are presented in figure 3.2. The behavioral results across conditions deviated no more than +/- 20% from the ideal behavior targeted by the design (50% for threshold, 100% for fully visible, 0% for scramble).

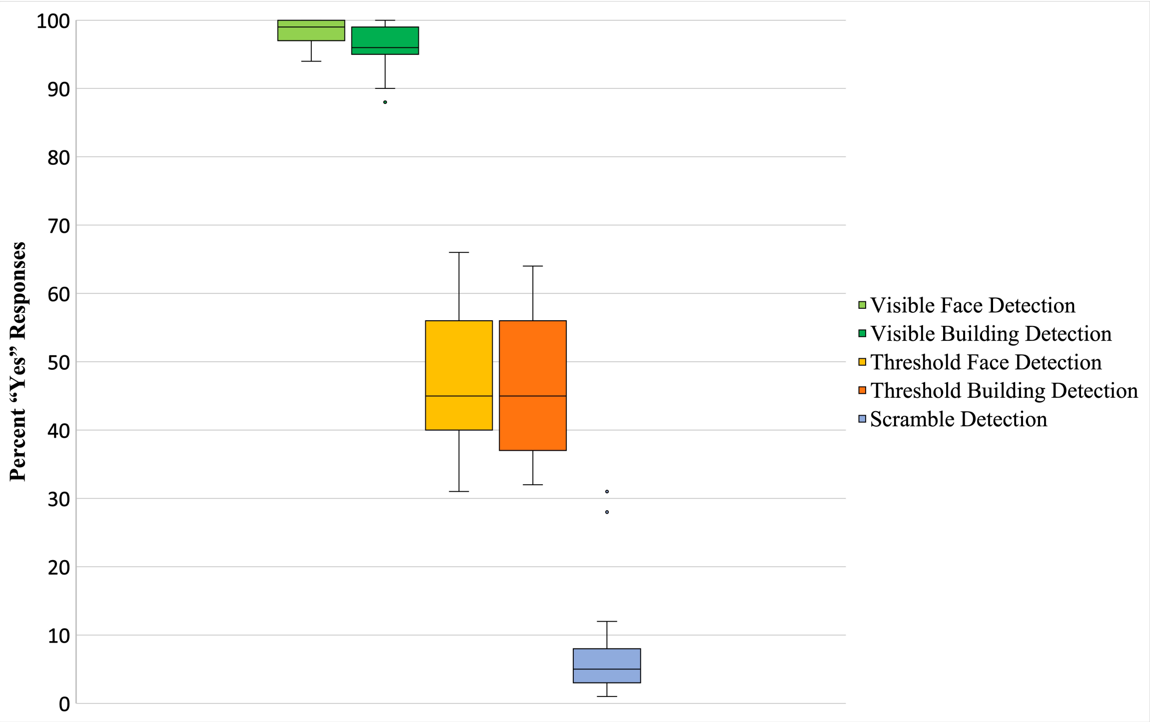

Figure 3.1 Detection Rate Across Conditions

Percentage of stimuli detected ("yes" response) in each condition for each participant.

Figure 3.2 Detection Rate Across Conditions by Subject

Percentage of stimuli detected ("yes" response) in each condition for each participant.

ERP Results

VAN and N170

The ERPs seen in figure 3.3 A-C below all show statistically significant difference in VAN amplitude. These significant differences in the time range of the VAN are present between fully visible faces and scramble, fully visible buildings and scramble, and seen threshold faces and unseen threshold faces. Seen threshold buildings and unseen threshold buildings, figure 3.3 D, did not yield statistically significant amplitude differences in the VAN time range (at least in the conservative mass univariate analyses), however, the topography in difference maps appears visually similar to the corresponding topography in visible buildings and scramble, and it is possible that more statistical power is needed to measure this small VAN in the threshold building condition.

It was challenging to distinguish the N170 from the VAN in the Analysis. In Figure 3.5, there is a clear time range of overlapping significance (170-270ms). In subsequent paragraphs, the N170 and VAN will be described together as the VAN, spanning the time window of both.

P3b

The ERPs seen in figure 3.3 A-C below all show statistically significant difference in P3 amplitude. Meaningful differences in the time range of the P3 are present between fully visible faces and scramble, fully visible buildings and scramble, and seen threshold faces and unseen threshold faces. Seen threshold buildings and unseen threshold buildings, figure 3.3 D, did not yield significant amplitude differences in the P3 time range, however, similar to the VAN in this condition, the topography seen in the difference maps looks visually similar to the difference topography for visible buildings and scramble.

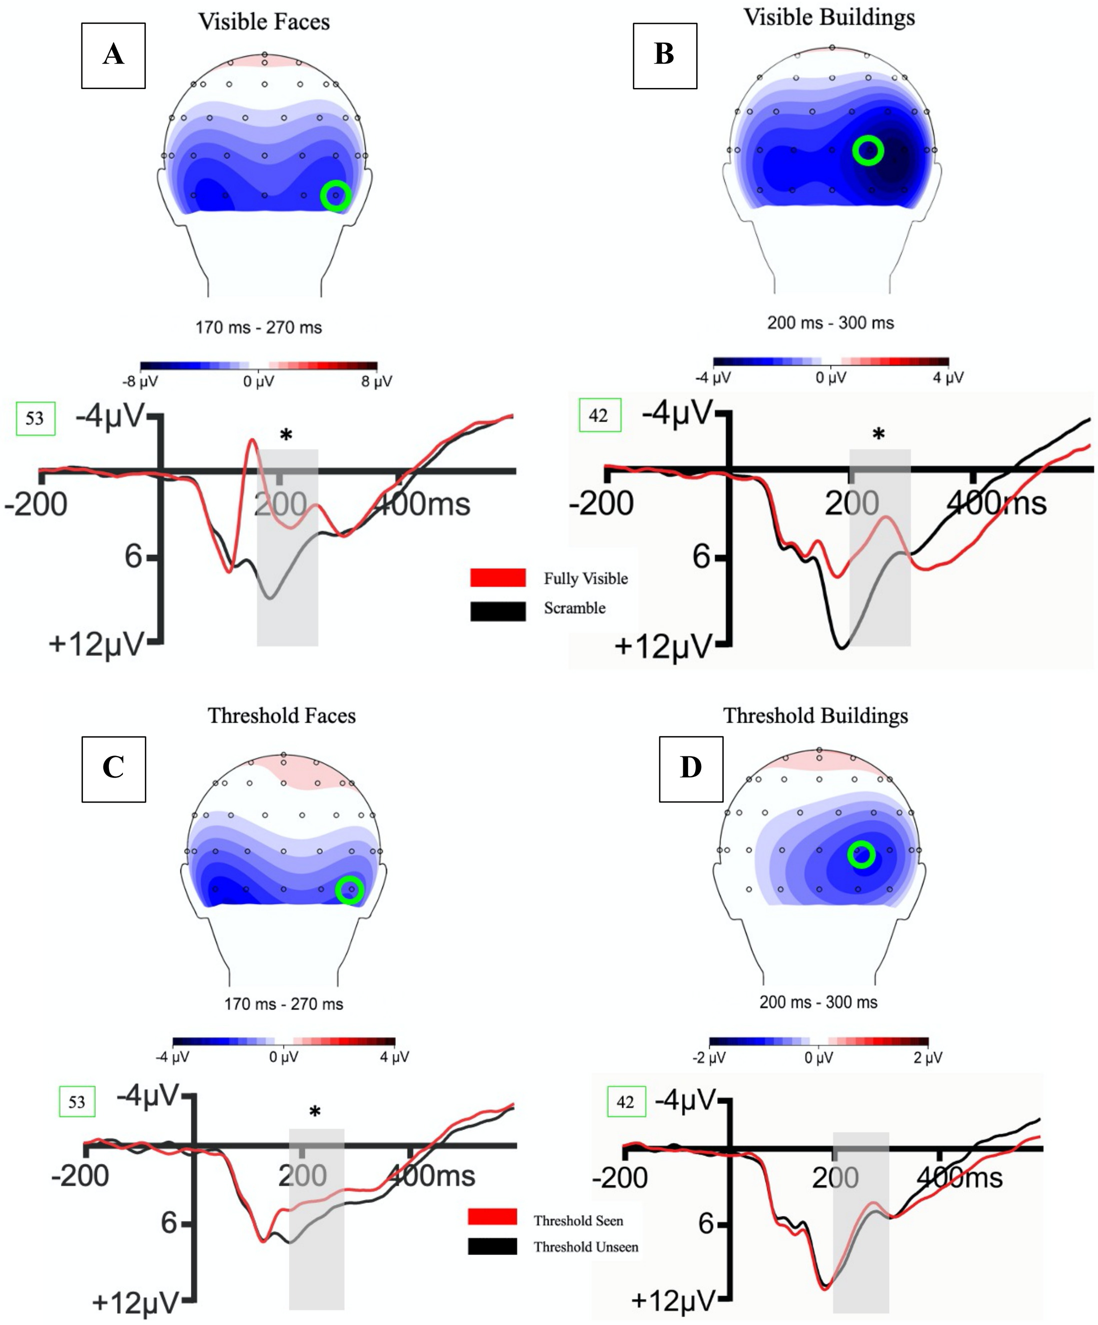

Figure 3.3 VAN in All Conditions: ERPs and Difference Maps

Grand averaged event-related potentials obtained A,C: Mean amplitude from 170-270ms after stimulus onset for fully visible faces versus scramble trials (A) and seen threshold faces versus unseen threshold faces (C) of electrode 53. B,D: Mean amplitude from 200-300ms after stimulus onset for fully visible buildings versus scramble trials (B) and seen threshold buildings versus unseen threshold buildings (D) of electrode 42. Gray boxed denote the time range depicted in the maps above. Stars denote a statistically significant difference in that time window ($p<0.05$). Difference maps showing the mean amplitude difference between either fully visible and scramble (A,B) or threshold seen and threshold unseen (C,D). Green circles denote the electrode used in the ERPs below.

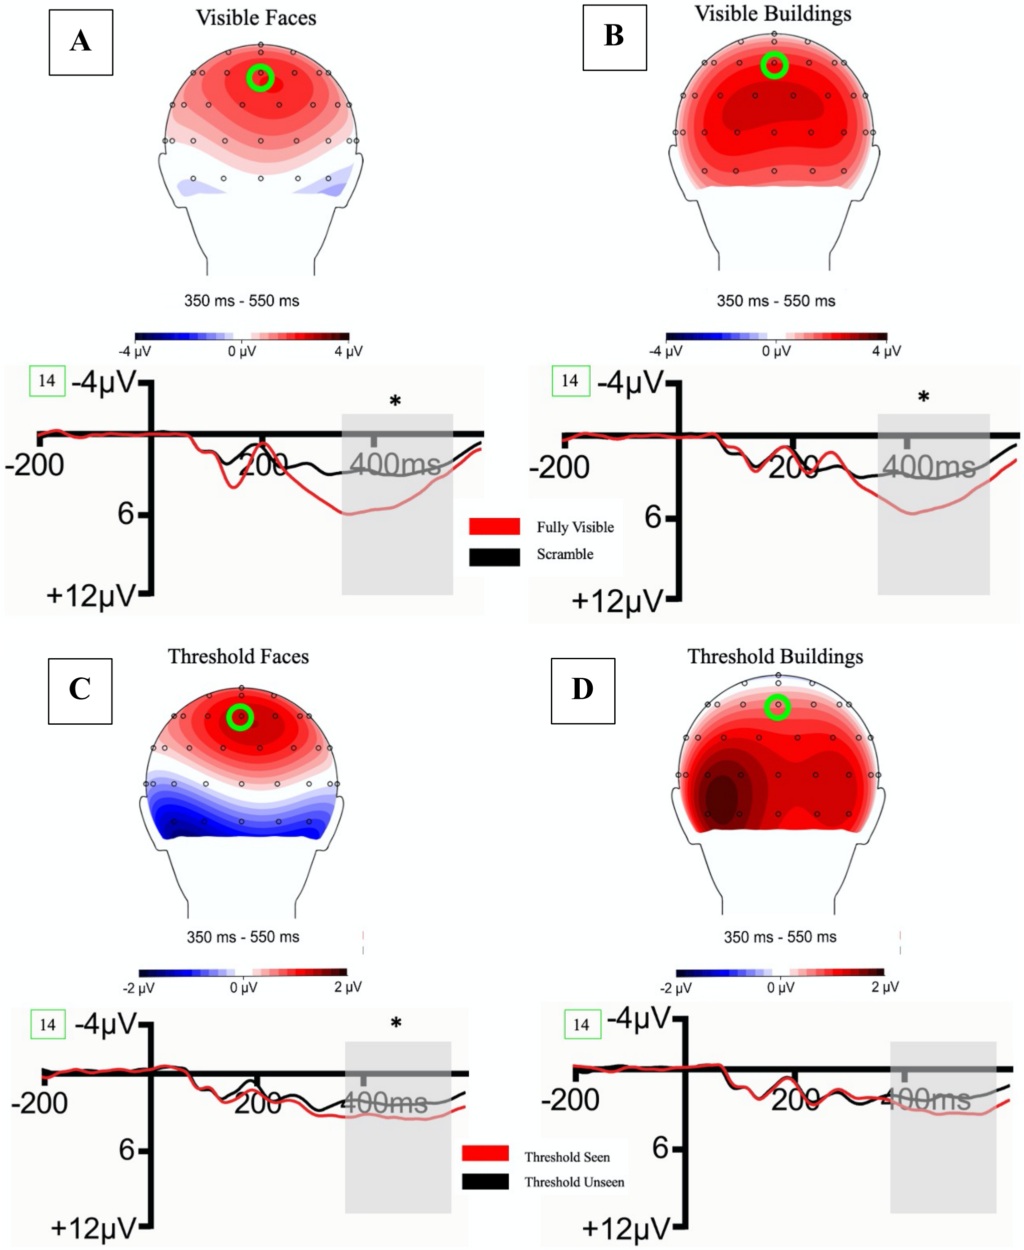

Figure 3.4 P3 in All Conditions: ERPs and Difference Maps

Grand Averaged event-related potentials obtained A,C: Mean amplitude from 350-550ms after stimulus onset for fully visible faces versus scramble trials (A) and seen threshold faces versus unseen threshold faces (C) of electrode 14. B,D: Mean amplitude from 350-550ms after stimulus onset for fully visible buildings versus scramble trials (B) and seen threshold buildings versus unseen threshold buildings (D) of electrode 14. Gray boxed denote the time range depicted in the maps above. Stars denote a statistically significant difference in that time window ($p<0.05$). Difference maps showing the mean amplitude difference between either fully visible and scramble (A,B) or threshold seen and threshold unseen (C,D). Green circles denote the electrode used in the ERPs below.

Mass Univariate Results

The plots below display the results of the mass univariate analysis, including differences between highly visible faces/buildings and unseen scrambles as well as between seen threshold faces/buildings and unseen threshold faces/buildings. Figure 3.5 and 3.6 show a significant difference (more negative voltages for seen versus unseen) in posterior electrodes around 200ms between highly visible stimuli and scrambled control stimuli for both faces and buildings, consistent with previous studies reporting a VAN.

Similarly, Figure 3.7 shows a pattern of neural differences between seen threshold faces and unseen threshold faces. Figures 3.5 and 3.6 show a significant difference in positivity in parietal-central regions after 300ms between highly visible stimuli and unseen scrambles for faces and buildings consistent with a P3b. Figure 3.7 also shows a similar pattern for seen threshold faces and unseen threshold faces. Figure 3.8 shows no significant differences at any time point between seen threshold buildings and unseen threshold buildings.

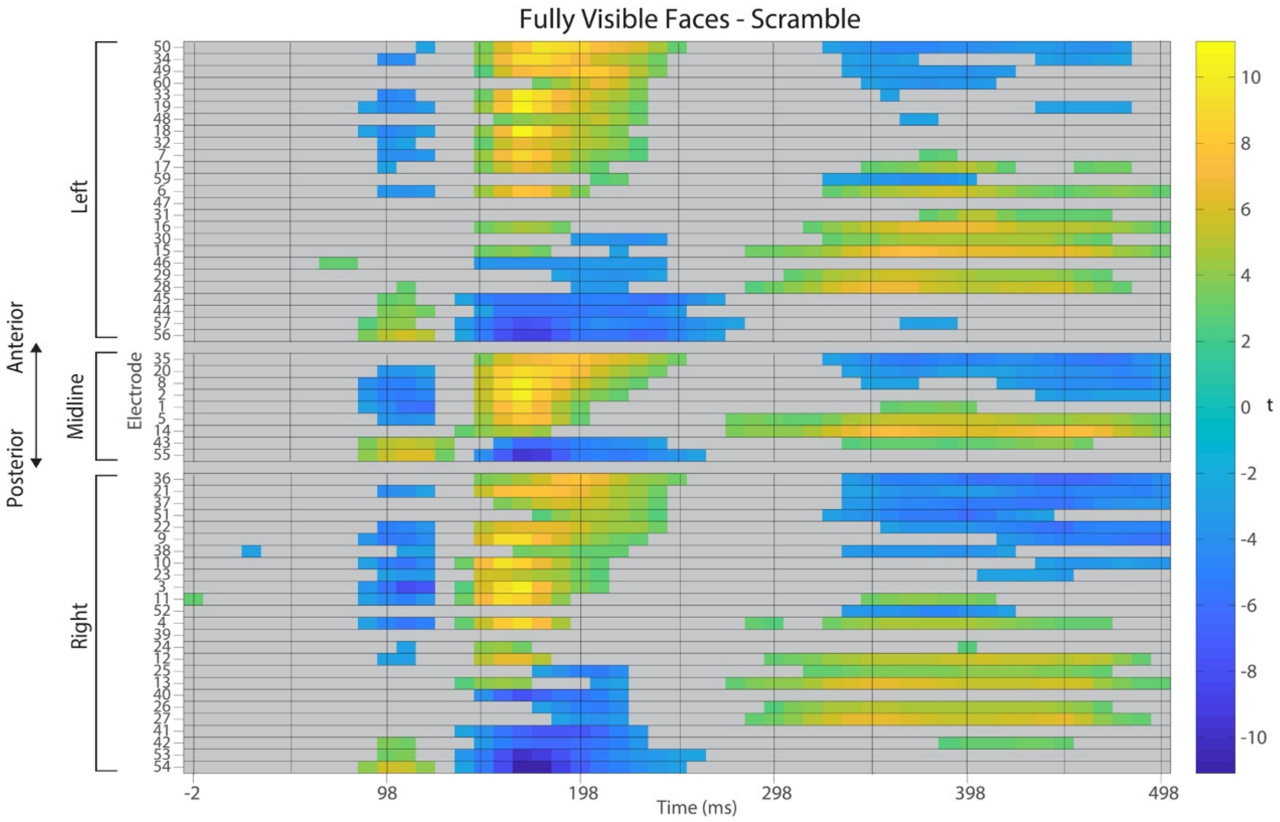

Figure 3.5 MUA: Fully Visible Faces Minus Scramble
Mass Univariate plot displaying statistically significant ($p < 0.05$) time points across relevant time window (0-500ms following stimulus onset).

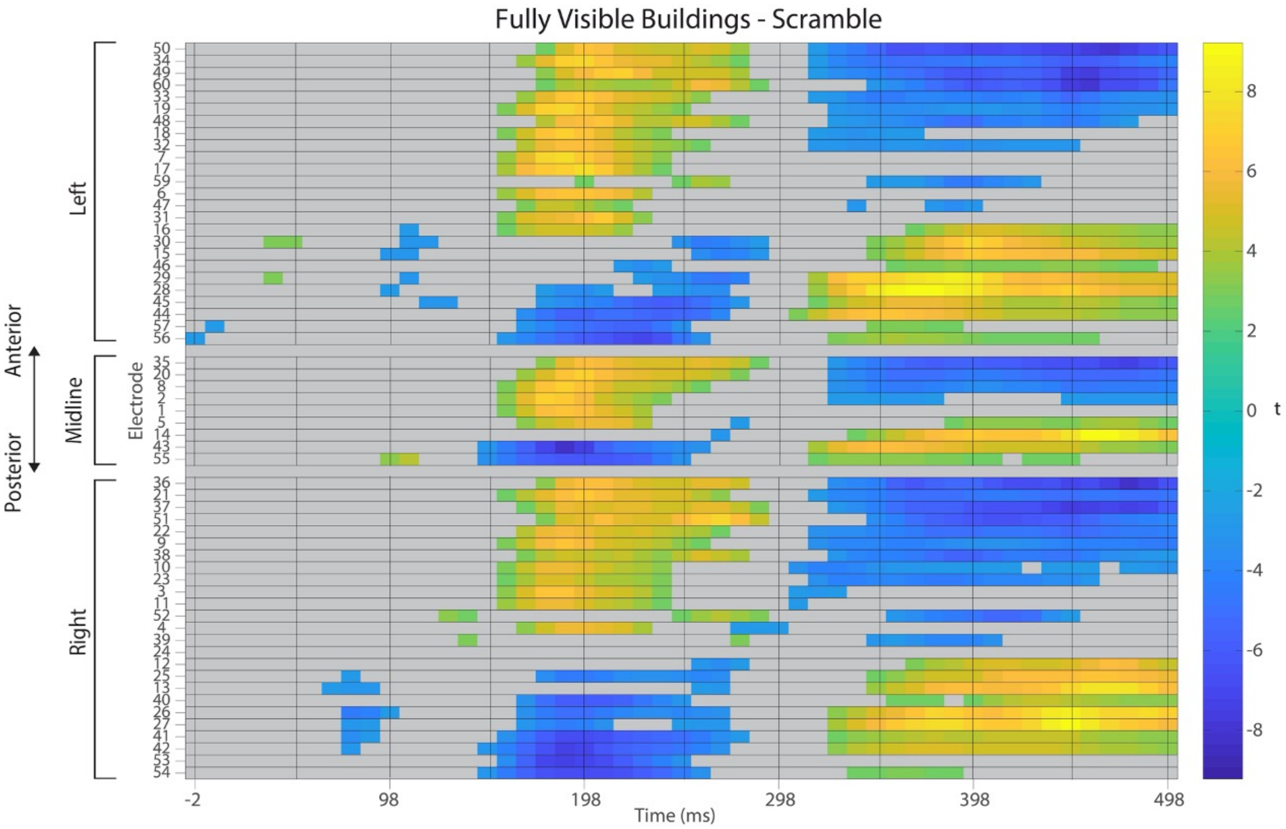

Figure 3.6 MUA: Fully Visible Buildings Minus Scramble
Mass Univariate plot displaying statistically significant ($p < 0.05$) time points across relevant time window (0-500ms following stimulus onset).

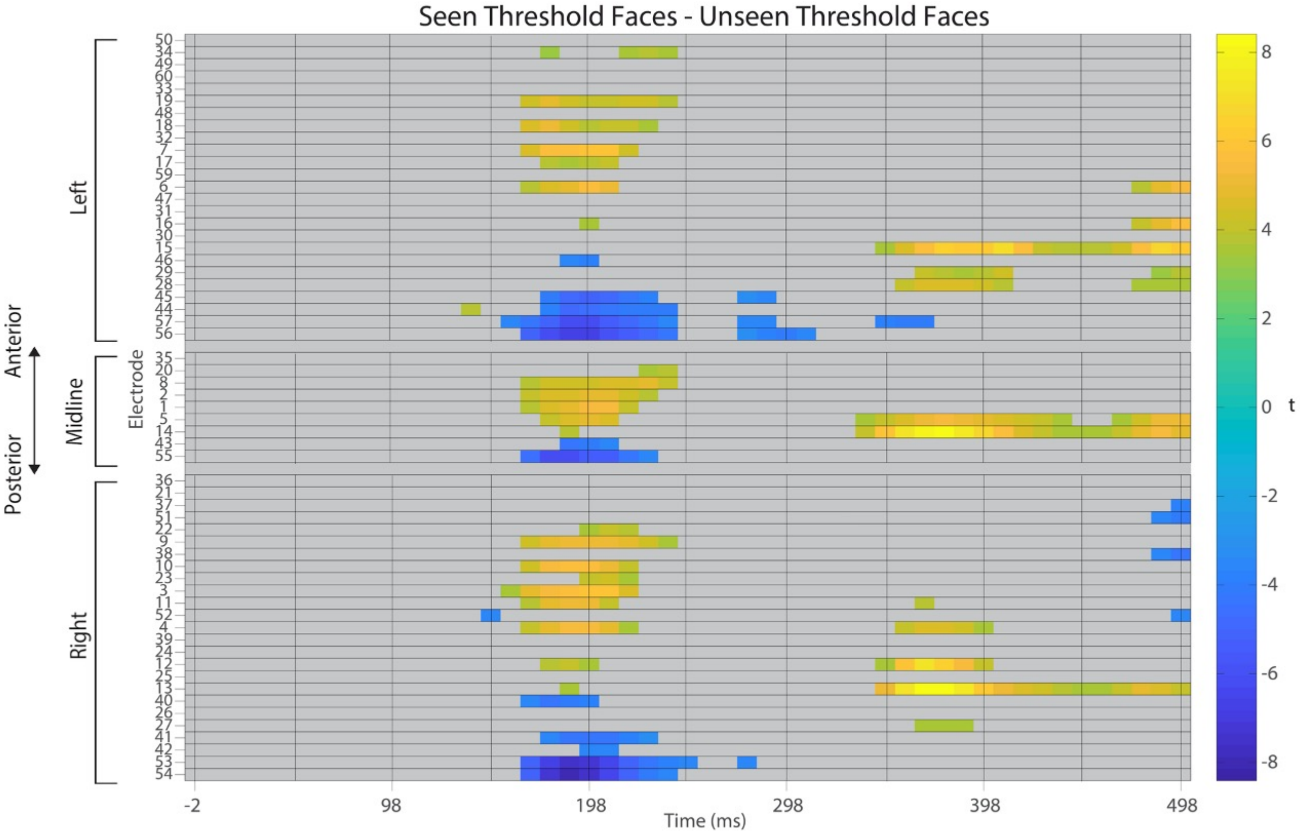

Figure 3.7 MUA: Seen Threshold Faces Minus Unseen
Mass Univariate plot displaying statistically significant ($p < 0.05$) time points across relevant time window (0-500ms following stimulus onset).

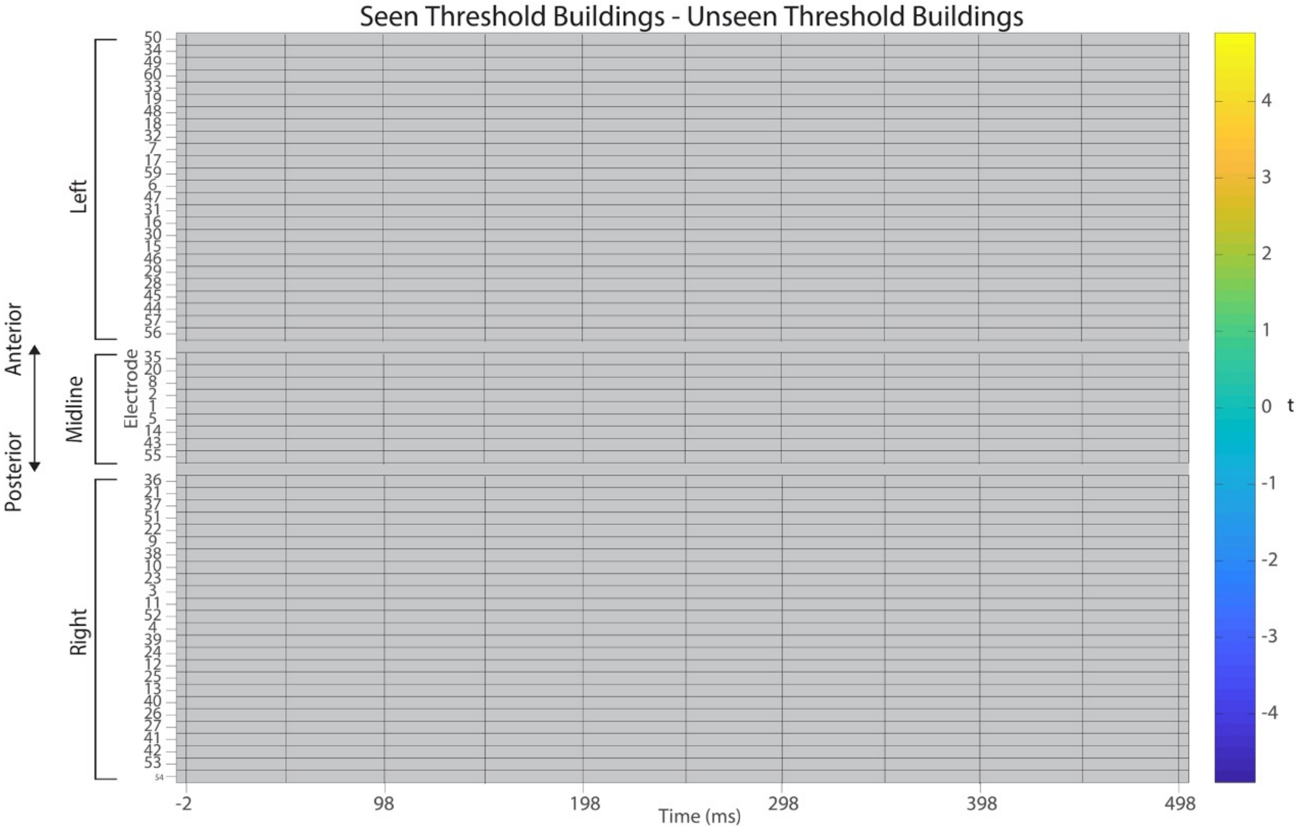

Figure 3.8 MUA: Seen Threshold Buildings Minus Unseen
Mass Univariate plot displaying statistically significant ($p < 0.05$) time points across relevant time window (0-500ms following stimulus onset).

Discussion

Behavioral Results

The behavioral results aligned well with performance measures targeted in the design of the experiment. The success of this paradigm to measure each individual subject's pattern perception thresholds and individually tailor and adjust the stimuli used in the main experiment to maintain near 50% detection of threshold stimuli across many trials is a compelling methodology for use in future studies, particularly those looking to measure neural differences based on perceptual awareness of physically constant stimuli.

ERP Results

As expected, the results of this study show significant differences in amplitude of the VAN and P3 between scrambles and highly visible faces, highly visible buildings, and threshold buildings. The contrast between highly visible stimuli and scrambled stimuli provide insight into the full set of potential neural differences in seeing a face or building and not, however it likely overestimates the neural differences as they include physical differences in the stimuli. As seen in figures 3.5 and 3.6, there is a clear increase in statistically significant effects compared to the threshold figures (figures 3.7 and 3.8).

The previous study in this lab, mentioned earlier, used very similar stimuli and found very similar neural differences (VAN), except for the absence of significant P3b differences. There was a P3b in response to both faces and scrambles (as well as scenes) in the report condition, but no significant P3b difference. It is likely that this may have to do with probability of the stimuli. A factor that leads to a P3b difference is probability in the stimulus presentation, with less frequent stimuli leading to an increased P3b amplitude. This experiment could be adapted in the future to present the stimuli with equal probability, potentially eliminating the P3b difference while maintaining the VAN.

The threshold condition in this experiment was used to narrow down which of the neural differences found in the visible versus scramble contrast might actually be related to conscious vision, by keeping the physical stimulus constant and only comparing seen and unseen. The differences seen for the threshold stimuli (figures 3.7 and 3.8) are significantly smaller and less extensive, meaning a majority of the differences in visible compared to scramble are due to physical stimulus differences.

The ERP results show a significant increase in VAN amplitude between visible and scrambled conditions as well as between seen and unseen faces, showing it to likely be a reliable index of visual awareness. Additionally, the results of the VAN demonstrate visual differences in topographical distribution on the scalp.

Topography Differences

The difference maps depicted in figure 3.3 display topographic differences between the face and building conditions. In face conditions, the difference of the VAN can be seen bilaterally over the occipital-temporal electrodes. In the building conditions, the difference can be seen occipitally with a subtle bias towards the right hemisphere. This finding poses the possibility of the VAN being category specific, or impacted by what one is consciously seeing.

Studies of auditory awareness negativity (AAN) suggest that auditory negativities are generated in specific cortices dependent on modality, rather than general regions (Eklund et al., 2019; Eklund et al., 2020). However, no previous study has explored whether VAN topography changes within a sensory modality for perceptual experiences with different categories of stimuli. The findings in the current study could offer a novel understanding of a category-dependent VAN, providing insight into the interaction between the VAN and category-specific regions.

A possible explanation could be the differences in the presentation or physical characteristics of the stimuli themselves. In figure 2.1, it is clear that the face stimuli extend from the center to the vertical length of the frame to the top and bottom of the frame (vertically), while the buildings tend to extend horizontally from the center of the screen to the right and left portions of the frame (horizontally). It is possible this could have impacted the variation in ERP results for the P3 and the VAN between stimuli. To combat this possibility, a future study could repeat the experiment, standardizing the spatial presentation of the stimuli by displaying the images behind a pre-defined circular border to mitigate spatial differences between buildings and faces, or through exploration of other visual categories of stimuli (letters, objects, shapes, color, motion).

Future Directions

The lack of a significant difference between seen and unseen threshold buildings for both the P3 and the VAN, as well as the visual similarity in scalp topography to the significant difference between fully visible buildings and scramble suggest that this effect was possibly not shown statistically as a result of the sample size.

The next step in understanding these conclusions is to replicate the study with more balanced probabilities in an effort to reduce or eliminate the differential P3b in threshold conditions. In the current experiment, seen threshold faces represented approximately 12.5% of all presented stimuli, whereas scrambles represented approximately 25% of all presented stimuli, which could explain the small difference in P3b amplitude. To investigate this further, the study could be replicated using only threshold stimuli, instead of including fully visible and scrambles, resulting in equal probabilities across conditions.

Another possibility could be to include variable ranges of stimuli targeted towards middle-range detection of the stimuli, such as presenting some that are perceived 50% of the time, some perceived 75%, and some perceived 25% of the time. Based on the findings of this study, it would likely result in an amplitude increase/decrease in the VAN, depending on the detection rates of these stimuli.

A final possibility could be implementing a no-report component to the experiment to evaluate the impact on presence and differences in the P3b, although this poses a difficult challenge in maintaining consistent detection rates throughout the length of the study.

Early vs. Late

The findings of this study align most closely with the outline of RPT, providing further evidence for the VAN, an early component, as a reliable indicator of awareness. These conclusions will likely prove more robust in the context of an expanded study designed to eliminate the late P3 difference through balancing probabilities, increasing the sample size, and eliminating physical differences in the stimuli.

References

American Psychological Association. (2020). Absolute Threshold. In *American Psychological Association*. Retrieved December 7, 2020, from dictionary.apa.org

Auksztulewicz, R., Spitzer, B., & Blankenburg, F. (2012). Recurrent Neural Processing and Somatosensory Awareness. *Journal of Neuroscience*, *32*(3), 799–805. https://doi.org/10.1523/jneurosci.3974-11.2012

Bayerl, P. A. J., & Neumann, H. (2010). Recurrent processing in the dorsal pathway underlies the robust integration and segregation of motion patterns. *Journal of Vision*, *2*(7), 658–658. https://doi.org/10.1167/2.7.658

Bentin, S., Allison, T., Puce, A., Perez, E., & McCarthy, G. (1996). Electrophysiological Studies of Face Perception in Humans. *Journal of Cognitive Neuroscience*, *8*(6), 551–565. https://doi.org/10.1162/jocn.1996.8.6.551

Bi, J., & Ennis, D. M. (1996). Sensory Thresholds: Concepts and Methods. Journal of Sensory Studies, 13(2), 133–148. https://doi.org/10.1111/j.1745-459x.1998.tb00079.x

Breitmeyer, B. G. (1984). Visual masking: an integrative approach. Clarendon.

Chalmers, D. J. (1995). Facing Up to the Problem of Consciousness. *Journal Of Conscious Studies*, *2*(3), 200–219.

Cohen, M. A., Ortego, K., Kyroudis, A., & Pitts, M. (2020). Distinguishing the neural correlates of perceptual awareness and post-perceptual processing. *The Journal of Neuroscience*

Crick, F., & Koch, C. (1990). Towards a neurobiological theory of consciousness. *The Neurosciences*, *2*, 263–275.

Dehaene, S., & Changeux, J. (2011). Experimental and Theoretical Approaches to Conscious Processing. *Neuron, 70*(2), 200-227. doi:10.1016/j.neuron.2011.03.018

Doerig, A., Schurger, A., & Herzog, M. H. (2020). Hard criteria for empirical theories of consciousness. *Cognitive Neuroscience*, 1-22. doi:10.1080/17588928.2020.1772214

Eklund, R., Gerdfeldter, B., & Wiens, S. (2019). Effects of a Manual Response Requirement on Early and Late Correlates of Auditory Awareness. Frontiers in Psychology, 10. https://doi.org/10.3389/fpsyg.2019.02083

Eklund, R., Gerdfeldter, B., & Wiens, S. (2020). Is auditory awareness negativity confounded by performance? *Consciousness and Cognition*, *83*, 102954. https://doi.org/10.1016/j.concog.2020.102954

Epstein, R., & Kanwisher, N. (1998). The Parahippocampal Place Area: A Cortical Representation of the Local Visual Environment. NeuroImage, 7(4). https://doi.org/10.1016/s1053-8119(18)31174-1

Förster, J., Koivisto, M., & Revonsuo, A. (2020). ERP and MEG correlates of visual consciousness: The second decade. *Consciousness and Cognition, 80*, 102917. doi:10.1016/j.concog.2020.102917

Glover, G. H. (2011). Overview of Functional Magnetic Resonance Imaging. *Neurosurgery Clinics of North America*, *22*(2), 133–139.

Harris, J. A., McMahon, A. R., & Woldorff, M. G. (2013). Disruption of visual awareness during the attentional blink is reflected by selective disruption of late-stage neural processing. *Journal of Cognitive Neuroscience*, *25*(11)

Hill, N. J., Gupta, D., Brunner, P., Gunduz, A., Adamo, M. A., Ritaccio, A., & Schalk, G. (2012). Recording Human Electrocorticographic (ECoG) Signals for Neuroscientific Research and Real-time Functional Cortical Mapping. Journal of Visualized Experiments : JoVE, (64). https://doi.org/10.3791/3993

Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The Fusiform Face Area: A Module in Human Extrastriate Cortex Specialized for Face Perception. The Journal of Neuroscience, 17(11), 4302–4311. https://doi.org/10.1523/jneurosci.17-11-04302.1997

Kim, C.-Y., & Blake, R. (2005). Psychophysical magic: rendering the visible 'invisible.' Trends in Cognitive Sciences, 9(8), 381–388. https://doi.org/10.1016/j.tics.2005.06.012

Mack, A., & Rock, I. (1998). *Inattentional blindness*. MIT Press.

Mashour, G. A., Roelfsema, P., Changeux, J., & Dehaene, S. (2020). Conscious Processing and the Global Neuronal Workspace Hypothesis. *Neuron, 105*(5), 776-798. doi:10.1016/j.neuron.2020.01.026

Owen, A. M. (2013). Detecting Consciousness: A Unique Role for Neuroimaging. *Annual Review of Psychology*, *64*(1), 109–133. https://doi.org/10.1146/annurev-psych-113011-143729

Pitts MA, Lutsyshyna LA, Hillyard SA. The relationship between attention and consciousness: an expanded taxonomy and implications for 'no-report' paradigms. Philos Trans R Soc Lond B Biol Sci. 2018 Sep 19;373(1755):20170348. doi: 10.1098/rstb.2017.0348. PMID: 30061462; PMCID: PMC6074089.

Polich, J. (2020). 50+ years of P300: Where are we now? Psychophysiology, 57(7). https://doi.org/10.1111/psyp.13616 Polich, J. (2020). 50+ years of P300: Where are we now? *Psychophysiology*, *57*(7). https://doi.org/10.1111/psyp.13616

Puce, A., & Hämäläinen, M. (2017). A Review of Issues Related to Data Acquisition and Analysis in EEG/MEG Studies. *Brain Sciences*, *7*(12), 58. doi:10.3390/brainsci7060058

Rodríguez V, Thompson R, Stokes M, Brett M, Alvarez I, Valdes-Sosa M, Duncan J. Absence of face-specific cortical activity in the complete absence of awareness: converging evidence from functional magnetic resonance imaging and event-related potentials. J Cogn Neurosci. 2012 Feb;24(2):396-415. doi: 10.1162/jocn_a_00137. Epub 2011 Sep 26. PMID: 21942763.

Shapley, R., & Hawken, M. J. (2011). Color in the Cortex: single- and double-opponent cells. *Vision Research*, *51*(7), 701–717. https://doi.org/10.1016/j.visres.2011.02.012

Schubert, T. M., Rothlein, D., Brothers, T., Coderre, E. L., Ledoux, K., Gordon, B., & Mccloskey, M. (2020). Lack of awareness despite complex visual processing: Evidence from event-related potentials in a case of selective metamorphopsia. *Proceedings of the National Academy of Sciences, 117*(27), 16055-16064. doi:10.1073/pnas.2000424117

Shuster, M. (1978). Approximate algorithms for fast optimal attitude computation. *Guidance and Control Conference*. https://doi.org/10.2514/6.1978-1249

Tsuchiya, N., Wilke, M., Frässle, S., & Lamme, V. A. (2015). No-Report Paradigms: Extracting the True Neural Correlates of Consciousness. *Trends in Cognitive Sciences, 19*(12), 757-770. doi:https://doi.org/10.1016/j.tics.2015.10.002

University of Washington Institute for Learning and Brain Sciences (Ed.). (2020). What is Magnetoencephalography (meg)? Retrieved April 15, 2021, from http://ilabs.washington.edu/what-magnetoencephalography-meg