At the Threshold of Consciousness: Investigating the Patterns of Neuronal Activity Events in an EEG No-Report Paradigm

---

A Thesis

Presented to

The Established Interdisciplinary Committee of Neuroscience

Reed College

In Partial Fulfillment

of the Requirements for the Degree

Bachelor of Arts

Maxwell Opitz

May 2025

Approved for the Committee (Neuroscience)

Michael Pitts

Acknowledgments

Mom and Dad, I could never thank you enough. Every year I develop a greater appreciation for the amount of time, energy, life and love that you have given me. Thank you for providing infinite support when I asked for it and when I didn’t. You are the most important and most formative people in my life. The only way I can hope to repay you is by becoming the version of myself you've made it possible for me to be.

Michael, thank you for the opportunities you’ve extended to me these past few semesters. I've been able to explore so much by getting involved in your lab. You've given me some incredible guidance and support this year. There were some areas of this research I knew very little about prior to starting, and it's hard to believe how far this has come–not just what I've completed, but how much I've learned. Last but not least, thank you for the shape tasting trip–that was a fantastic time.

Angelica, thank you for helping me out basically every day. If I didn't have you to watch and learn from, I would've been so much less assured throughout this process and I probably would have made some much more egregious mistakes. It’s been a pleasure working with you this year and I really look forward to this summer.

Erik Zornik and Mark Hinchliff, thank you for reading my thesis and providing an incredibly interesting 90 minutes of discussion.

Max Bennett, thank you for coding the programs we used this entire year, and basically making the whole experiment possible.

To my friends who volunteered for this thesis: thank you for making this process not only possible but pleasant. Coming in at odd hours to sacrifice your time and your brainwaves is a real commitment, and I'll forever be grateful.

Ellen, thank you for being an anchor amongst the turbulence of life and an endless source of support. Thank you for showing me that spreadsheets aren't actually that daunting (except for when they kinda are). Thank you for generally helping me not lose my mind on any number of occasions.

Grandpa Lyn, thank you for making so many things possible for me, and for being there for me at so many points in my life. This journey would not have been possible without you, and I’m so glad to have you in my life!

Grandma Leslie and Grandpa Chuck, thank you for being such incredible grandparents and figures in my life. You've provided so, so much for me and I've learned so much from you. Grandpa, you will always be a source of wisdom and guidance to me, and I'm so grateful for the time I had with you. I miss you, and I love you always.

There are definitely more family and friends I should include here. I was aiming for under a page of acknowlegements, so I think cutting it off at two is best. Nonetheless, to everyone else who helped and inspired me, truly, thank you. I appreciate you so much.

List of Abbreviations

NCC	Neural Correlate of Consciousness
GNW	Global Neuronal Workspace
V1	Primary Visual Cortex
EEG	Electroencephalography
ERP	Event-Related Potential
PFC	Prefrontal Cortex
VAN	Visual Awareness Negativity
PAN	Perceptual Awareness Negativity
P3b	Positivity @300 milliseconds (type b)
FcN2	Fronto-central Negativity @200 milliseconds
PAS	Perceptual Assessment Scale
OSF	Open Science Framework
2afc	2-Alternative Forced Choice
LFN	Late Frontal Negativity
LVAP	Late Visual Awareness Positivity
3C	Concurrent, Competitive, Cognitively Involved

Table of Contents

List of Tables

List of Figures

Abstract

A promising area in the ongoing effort to understand human consciousness is the pursuit of Neural Correlates of Consciousness (NCCs)–the underlying signals and mechanisms which actually give rise to the conscious experience. Recently, Cohen et al. (2024) studied EEG-recorded signal responses to visual stimuli of varying visibility, measuring how neural mechanisms behave above, below and at the very threshold of what is consciously perceptible. The study also used a no-report/report paradigm, in which EEG recordings were taken without the subject needing to actively report their perception of a stimulus. By utilizing a means to compare the signal amplitudes of unreported perception with a non-linear function of visibility, several signals could be investigated for their role in pure conscious perception indepedent of attention and sensory processing. This thesis followed up Cohen et al. (2024), simplifying the paradigm further by exchanging masked facial stimuli for simple pattern gratings of varying contrast against visual noise. EEG recordings from 10 pilots and 30 final experimental subjects were collected, alongside separately collected behavioral measures to determine the visibility of each stimulus contrast level. One signal of interest from the previous study, which closely matched the non-linear pattern of awareness in the report condition, was a fronto-central negativity (fcN2) from ~250-300ms. In this thesis, despite being the first study to attempt to directly isolate it, the signal did not appear; extensive literature review of the conditions which evoke this signal indicate a complex relationship to attentional perception and possible avenues for further experimental isolation. Also of interest was the earlier Visual Attention Negativity (VAN), a signal with a promising but disputed status as a marker of consciousness. Its patterns in this study tie it more closely than ever to the generation of perceptual awareness, but with unique characteristics that potentially inform the nature of the no-report

condition itself. Throughout the study, a late signal system consisting of a frontal negativity and bilateral posterior positivity were also discovered and investigated. These signals formed intriguing spatial and amplitude spatterns which could be theoretically significant.

PART 1: INTRODUCTION

1.1. Overview: The Study of Consciousness

One of the most exciting frontiers in neuroscience is the study of consciousness–not a single human experience, but rather the state of *experiencing* itself. How does a mind generate this fundamental awareness of itself, to not just *be* but *experience being*? Consciousness manifests as observation, sensation, attention, and every other subjective element which bridges the mind's gap between itself and the world. An ongoing principal effort in neuroscience has been to determine the corresponding components of the human brain responsible for generating these elements, and through them, consciousness as a whole. This includes building models of how consciousness arises, parsing which anatomical areas and cell types are responsible, and isolating the particular neuronal mechanisms which mechanistically create it (Koch et al., 2016). These mechanisms, called Neural Correlates of Consciousness (NCCs), are the specific circuitries and activity events that can be considered the minimum necessary to produce the conscious experience (Koch et al., 2016). As the last three decades of research have continued to discover and isolate which areas of the brain comprise full and content-specific NCCs, theories as to how consciousness broadly functions in the human brain have concurrently been developed and disputed (Mashour et al., 2020). While research continues to shape the idea of consciousness across various scholarly fields–from neurobiology to philosophy–various models have been synthesized which, at this time, offer promising insight into the question of what consciousness really *is*. As such, it is useful to begin with an overview of what is currently understood and what models offer the best syntheses of those findings.

1.2. How Consciousness is Modeled

1.2.1. The Phenomenal/Access Distinction, Shortcomings, and Alternatives

A longstanding theoretical distinction in consciousness research is between "phenomenal" and "access" consciousness (Block, 2005). Access, in this case, refers to how conscious information is available to various cognitive processes such as motor behavior, working memory and reporting–in other words, a conscious experience is *accessed* when it is understood or acted upon in any way. Proponents of the distinction believe this accessed experience is only a subsection of the phenomenon of consciousness itself, which some posit as a sovereign, pure subjective experience (sometimes called *qualia*) that hypothetically exists independent of how it is actually accessed and processed by the brain's cognitive functions. The root of this distinction is in whether one's access to (perception & memory of; ability to act upon) their consciousness is the determining factor in whether the experience itself is actually conscious or not. The importance and validity of this distinction has been continuously debated, as it is dubiously possible to experimentally test the existence of a conscious experience that cannot be accessed by the entity which possesses it. Even awareness of one's own consciousness implies a transitive "access," and if someone is entirely unaware of being conscious, what distinguishes it from unconsciousness–or from beings and objects that don't experience consciousness at all? This search to detach conscious phenomena from cognitive access, some argue, leads to definitional indistinction and imprecision (Naccache, 2018).

While there has been limited speculation on neuroanatomical functional distinctions between phenomenal and access consciousness ((Lamme, 2006, 2018, p. 201), the aforementioned lack of experimental possibility–in tandem with growing logical concern–has directed some models to view "access" as a given and to characterize delineations as mediated by subjective reportability (Naccache, 2018). This does not suggest everything within the conscious experience can be fully reported; rather, to the contrary, that accessible experience is much broader than what can be reported. Rather than the framing that "we can phenomenally experience more than we can access", it may be more accurate to say that "we can access more than we can report." This delineation utilizes the level (overall state, such as "attentive" or "drowsy") and content (type of information being perceived) of conscious processing as measures to determine the defining lines of what is necessary for a proper conscious state (Bachmann & Hudetz, 2014).

1.2.2. The Global Neuronal Workspace

One model which has been updated and sustained alongside contemporary developments in neuroscience is the *Global Neuronal Workspace* theory, which suggests that the availability of transmitted information across a vast "global" network of specific neurons determines the condition that is subjectively experienced as being conscious (Dehaene & Naccache, 2001). GNW Theory was proposed by Dehaene et. al. (1998) based on the "global workspace" framework originally proposed by Baars (1988). The GNW theory adopts Baars' tenets alongside newer research into substantiated neuroanatomical and neurobiological mechanisms. It takes into account the valuable insights offered by other models such as feedback loops (Recurrent Processing Theory) and the significance of the prefrontal cortex (Higher-Order Theory), offering one of the more compelling semi-unified models of conscious experience generation based on available research (Mashour et al., 2020).

The central idea of the GNW theory involves the existence of a defined "workspace" consisting of millions of neurons distributed widely across the brain, each acting in coordination to compute and communicate signals.

Information doesn’t begin in the workspace–rather, it is first processed in specialized local cortices mainly pertaining to sensory input, such as the Primary Visual Cortex (V1). After this local processing, there is a sudden and coordinated ignition of activity in which the information is swept from the bottom-up to the higher-order areas of the brain—the cortical zones that compose the bulk of the upper region—in which the workspace operates (Figure 1). Within a few hundred milliseconds, these working neurons go from receivers to broadcasters, retransmitting the signals back to local processing areas, establishing a tight bidirectional connection. This feedback from lower- to higher-order cortices and back down again creates signaling loops which remain recurrent, and it is in the establishment of these continual loops, the theory posits, that consciousness is generated and maintained. If either direction of the loop is disrupted, so is consciousness.

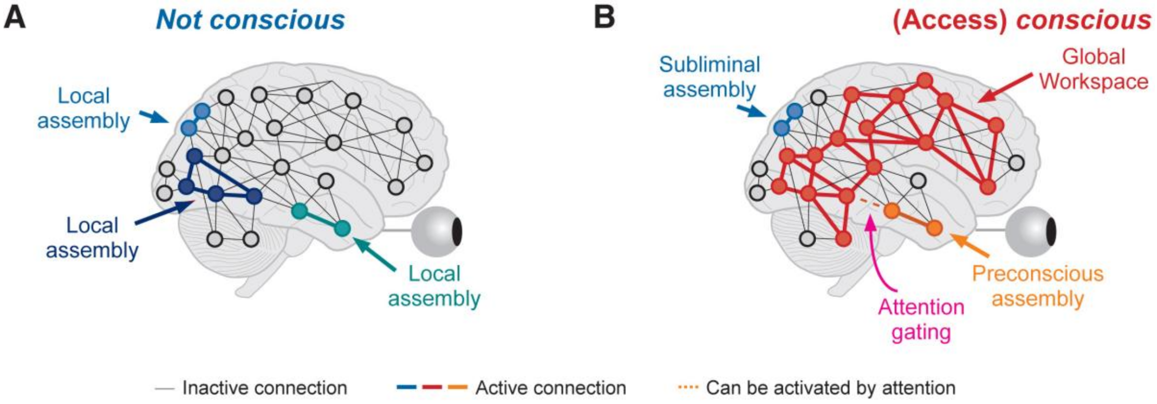

Figure 1. A diagram of the Global Workspace.
The role of the Global Neuronal Workspace in conscious processing from (Storm et al., 2024).

1.3. Biological Overview

1.3.1. Neurotransmission

Understanding the neurobiology and cellular specializations involved in consciousness-specific signaling can offer insight into how these processes function; furthermore, identifying which molecular components lend themselves to consciousness can corroborate the importance of cortices and signaling mechanisms in which those components operate. After all, the brain is a biological organ, and examining the biological units which construct the mechanisms of consciousness may be one of the most valuable tools for understanding it. As with any neurobiological process, the excitation and suppression of the signals which generate consciousness are mediated by neurotransmitters and their associated receptors. The excitatory neurotransmitter glutamate seems to be a primary driver; AMPA glutamate receptors are the largest facilitators of the initial concurrent "feedforward" ignition from lower processors up to the centralized higher-order workspace, while NMDA receptors facilitate the subsequent top-down broadcasting the workspace performs (Self et al., 2012). This corroborates the pharmacodynamic findings that manipulation of glutamate pathways (e.g. how much transmitter is binding to NMDA vs AMPA) induces complex changes to sensory experience, and furthermore that NMDA receptor availability modulation can significantly alter or impair conscious perception (Berkovitch et al., 2017). This connection furthermore supports the GNW model's notion that consciousness is not stable until its signals achieve a complete bidirectional feedback loop.

1.3.2. Cellular Specializations

Key to workspace signaling are pyramidal neurons, which possess specialized axons capable of long-range reception and transmission to generate and sustain these aforementioned tight feedback loops. These specializations have been observed abundantly in the 2nd and 3rd cortical layers, offering insight into the anatomical geography of the workspace (Mashour et al., 2020). Also robustly associated with consciousness are the pyramidal and Von Economo neurons appearing densely in the deeper fifth cortical layer in the frontal lobe, which serves largely to project signals between cortical and subcortical tissues (Butti et al., 2013; Yu et al., 2008).

Both thick-tufted and thin-tufted pyramidal neurons utilize feedback integration loops with subcortical and cortical interconnected regions, respectively (note that “tuft” thickness refers to the density of branch-structures on the dendrite, which enables input and output of synaptic transmission). Because of this, it is thought that the thin-tufted variants in particular may play a role in facilitating long-distance loops between the higher-order cortical zones throughout which the workspace neurons are distributed (K. D. Harris & Shepherd, 2015; Larkum, 2013). In the densely interconnected supragranular layers of pyramidal neurons, specialized excitatory feedback events called *avalanches* are transmitted in connection with receptive matrix cells in local lower areas such as the thalamus via specialized apical dendrites (K. D. Harris & Shepherd, 2015; Sakata & Harris, 2009; Zhang et al., 2014).

The role of such subcortical zones as the thalamus is also cellularly significant. Robust transmission behavior of cortical neurons certainly indicates wide-reaching global distribution functionality, and recent evidence of a "gating" role conducted by the thalamofrontal loop corroborates the intralaminar and medial nuclei specifically (Fang et al., 2025). This further lends evidence to the 2nd, 3rd, and 5th layers of the prefrontal cortex as an imperitative destination for integrating information into consciousness.

1.4. Isolating Signals

1.4.1. How Signals are Found–EEG and ERPs

It's becoming clear through this overview that modeling the generation of consciousness is dependent largely on understanding what signals are being generated at what times by particular components of the brain. One of the foremost tools for measuring brain signaling in response to controlled stimulus is electroencephalography (EEG), in which electrodes are non-invasively attached to a participant's scalp to measure the voltage activity, millisecond-by-millisecond, of the electrical impulses that compose neurological signaling across the brain's geography (Britton et al., 2016). As these electrodes are on the surface of the scalp, the signals they detect are generally the product of neurons near the surface of the brain–the aforementioned cortical pyramidal neurons of interest. When the previously described neurotransmission processes of excitation and inhibition occur, these neurons achieve postsynaptic potentials; the immediate consequence of this chemical transmission is an electrical voltage event. Often great swathes of neurons potentiate in synchrony, leading to measurable coordinated activity events in particular areas (Britton et al., 2016).

With EEG, these events can be measured in real-time and later analyzed at precise time signatures to determine how they correspond to experimental variables, such as the presentation of a stimulus to the participant. An electrophysiological activity which can be isolated as the direct result of a particular event, such as a sensory perception or motor function, is called an event-related potential (ERP). An EEG for even a single participant yields abundant data, and preprocessing can reveal promising patterns, but the results become even more informative when brainwave information is aggregated across multiple electrodes and a full set of participants. Noise is reduced and the preserved waveforms become more clear; the consistent presence of a signal at a particular time signature which corresponds to an event indicates that the event may be causing the signal. It is through the discovery of ERPs that much of the timing and spacing of the brain's signaling has been illuminated.

1.4.2. Early Processing

EEG research has corroborated a central tenet of the GNW theory–that before being computed into conscious perception, sensory information is processed by recurrent looping in lower-order, local sensory processing areas which don't correspond to awareness. This processing is understood to occur, on average, for the first 200ms after a stimulus (such as a visual shape) is presented (Marti et al., 2015; Marti & Dehaene, 2017; Sergent et al., 2005). In fact, it's been demonstrated that these signatures of early processing are identical in trials whether or not a participant reported actually consciously perceiving anything, affirming that consciousness of a sensation is not processed in its corresponding sensory zone (e.g., the visual cortex processes visual information but does not create the conscious experience of having vision). If a participant is distracted, the duration of this early processing can increase well beyond 200ms (Marti et al., 2015), and certain signals related to consciousness may ignite within this window, but it serves as a general functional distinction point. Thus, signals which persist after at least 200ms–which diverge depending on whether the stimulus was perceived–are of the greatest interest in locating which signals generate that perception.

1.4.3. Neuroanatomical Overview–Areas of Interest

The GNW theory posits that no area involved in the workspace is redundant, with each playing a unique role in this process of broadcasting and transmission (Mashour et al., 2020). This theory and others have widely disputed the importance of various areas of the brain in sustaining the full NCC, though research via EEG and other methods has offered varying levels of certainty as to the involvement of certain components.

Subcortical Research.

Neural tissues below the surface-level cerebral cortex are more difficult to study through EEG, but lesion studies and a variety of modulatory techniques have offered insight into which structures may or may not play a role in consciousness. Many subcortical structures, such as the brainstem and its associated baseline mechanisms, are critical for survival and wakefullness but can be considered, in that sense, background requisites-they do not generate consciousness directly but provide conditions necessary for it. Even when the activity of structures such as the cerebellum and basal ganglia are dormant, thus preventing wakefulness, patients have reported being conscious under certain conditions (Parvizi & Damasio, 2001).

The thalamus is a rapidly developing area of subcortical research. As mentioned in the cellular overview, it possesses receptive matrix cells for sustaining thalamocortical loops over the long distances with the cortical neural network. While in studies the correlation between thalamic function and consciousness has varied significantly (Alexander et al., 1986), it had been long suspected that top-down communication to this area and other subcortical tissues plays some necessary (but not directly generative) role. Recently, it has been discovered using intracranial EEG in mice that the thalamofrontal loop plays a "gating" role, maintaining or blocking the transmission path for each bit of information traveling from primary sensory areas to the cortical network neurons (Fang et al., 2025). This mediation is a specific, compelling instance of an actual mechanistic operation that occurs during the activity sweeps between subcortical and cortical zones.

Cortical Research.

As previously mentioned, there is growing evidence that certain layers of the cerebral cortex, particularly in the frontal lobe, play a significant role in neurotransmission signaling related to consciousness. The prefrontal cortex, given its aforementioned high density of critical broadcasting neurons, is understood to play a key role in generating consciousness. While some models, such as Higher-Order Theory, posit that the Prefrontal Cortex (PFC) is the exclusive area of consciousness generation, a model which accounts for the role of inter-cortical and long-distance feedback signaling is more cohesive with current research. So, while the PFC plays a necessary role, its circuitry is not exclusively capable of generating consciousness. What remains more complicated are specific areas of the prefrontal cortex. While the inferior parietal and dorsolateral prefrontal areas were originally theorized to be central workspace regions (Mashour et al., 2020), revelations in the role of cognitive attention have called certain signals into question.

1.4.4. Signals of Interest

The VAN.

One of the earliest potential signals associated with conscious perception is the Visual Awareness Negativity (VAN), observed in the 100-300ms time range (note it is considered the AAN or SAN in auditory and somatosensory perception research, respectively, but it is more often studied using visual stimuli) (Dembski et al., 2021). The VAN occurs in posterior regions and robustly exists as an ERP associated with conscious visual perception, though its exact role has been stipulated. Those in support of a phenomenal vs. access consciousness distinction regard that it may be more related to the former (Eklund & Wiens, 2018), and the question remains of whether it correlates specifically with perceptual awareness or with something else–perhaps sensory strength or a certain type of visual attention (Cohen et al., 2023; Dembski et al., 2021). The existence of a general marker of perceptual consciousness at this posterior point (perhaps most aptly labeled the "Perceptual Awareness Negativity (PAN)" (Dembski et al., 2021)) is promising as a means of understanding the initial neural events that occur upon stimulus perception. Further research into how exactly this negativity behaves under varying conditions will likely prove illuminating.

The parietal cortex and the P3b.

Perhaps the most major–and controversial–signal of discussion in NCC research is the P3b, a late (300ms and beyond) fronto-parietal ERP which has been robustly correlated with auditory and visual stimulus detection. While the P3b was considered an NCC for some time, subsequent revelations have indicated it may be more associated with the function of reporting and task execution, making it a marker more of cognitive activity which depends on conscious content than of consciousness itself. Consider that when one presses a button to indicate they have seen something, the resulting neural activity could be indexing the seeing of the thing *or* the recognition and decision-making pathways that lead to pressing the button. Indeed, various studies have proven the P3b is only present when the participant regards a stimulus as relevant to their current task and is anticipating its presentation in an attentive manner (Cohen et al., 2020; Pitts et al., 2012; Pitts, Metzler, et al., 2014; Pitts, Padwal, et al., 2014; Schelonka et al., 2017; Shafto & Pitts, 2015).

Despite this robust rebuke of the P3 as a proper NCC, controversy remains. On a definitional level, the role cognitive attention necessarily plays in conscious must be considered, as well as whether it can be full divorced from the generation of a conscious experience at all. Indeed, if a stimulus is completely irrelevant to an observer, and unattended entirely by their attentive process, can it be considered consciously perceived? (Pitts et al., 2018). Putting a fine point on what this signal does and does not index may require further reconciliation between phenomenal, accessed and reported experience.

1.5. The Previous Study

1.5.1. Purpose, Paradigm and Findings

Cohen et. al. (2024) gathered electroencephalography (EEG) data to search for neural correlates of perceptual awareness. In this study, participants were presented with a face which was immediately followed by a “mask”—an obfuscating shape–with varying time in between (some tens of milliseconds). Altering how rapidly the mask appeared would toggle whether the participants were aware that a face had been presented before it. By controlling whether visual awareness occurred, the researchers intended to measure and discover ERPs associated with conscious perception of a stimulus.

The task paradigm was structured with an emphasis on distinguishing the activity of a perceptual awareness NCC from neighboring excitatory activity induced by early sensory recognition or by later reporting and attention-induced activity. The first element of this paradigm was to manipulate the timing of the masks such that sensory strength would increase linearly and awareness would increase non-linearly. The central basis of this idea is that awareness is a binary condition–either one perceives a thing or they do not; a stimulus is either invisible or visible. Theoretically, the singular jump between these unchanging conditions would be the point where a stimulus crosses a threshold from not perceptible to perceptible.

By using five different lengths of time between the face and the mask, this threshold point was established, wherein the face was almost never seen at the two shortest intervals (17 & 33ms), almost always seen at the two longest (67 &

83ms) and seen 50% of the time for the moderate duration (50ms). As such, activity associated with sensory strength and awareness increased on different curves–the former linear and the latter in a bifurcated S-curve (Figure 1); thereby, plotting ERP amplitudes evoked by each variation would indicate whether it correlated specifically with perceptual awareness.

The second element of the paradigm was to include a no-report condition (Frässle et al., 2014; Kapoor et al., 2022; Pitts et al., 2012) in tandem with the report trials. This meant that alongside "reporting" trials (i.e., the participant pressed a button to indicate they had noticed the presence of a stimulus), each participant additionally underwent trials in which they did not take any action at the time of perception. Instead, extrapolating from the report condition, physiological measures were utilized to infer when the participant had noticed the stimulus. When isolating a particular ERP from preceding and subsequent activity events, a no-report segment can be greatly beneficial for clarifying the data. In this particular study and others (Cohen et al., 2020; Hatamimajoumerd et al., 2022), certain signals of interest–such as the P3b and temporal generalization signals–which were discovered in the report trials were absent in the no-report condition, further supporting the growing consensus that such signals correspond moreso with the task of reporting and attentional focus than with conscious perception of the stimulus itself.

Through this split-condition and interval-variation paradigm, the study discovered an ERP signal of interest which appeared to be significantly linked to perceptual awareness independent of reporting and sensory activity–the *fronto-central N2*, present from ~250-300ms. Interestingly, the N2 wave appeared only in the no-report condition. This implies the interference effect of the report-associated event activity (namely the P3) was sufficient to drown out and obscure this signal in previous report-reliant paradigms–a hypothesis supported by post-hoc exploratory analysis of the study. The fcN2 was also consistently coupled with posterior positive potentials which were larger on the right hemisphere–this was universal across all observed instances of the N2, even in other studies in which an N2-esque ERP was only later discovered in retroactive analysis (Sergent et al., 2021). Furthermore, though it is negative and within the same ~150-300ms time frame as the Visual Awareness Negativity, the N2 is notably more anterior (Förster et al., 2020). It appears the fcN2 is conserved for auditory awareness as well, but the positive posterior event coinciding with it switches from occipital to temporal. These anomalies, as well as the presence of the N2 itself, warrant further research into its exact nature.

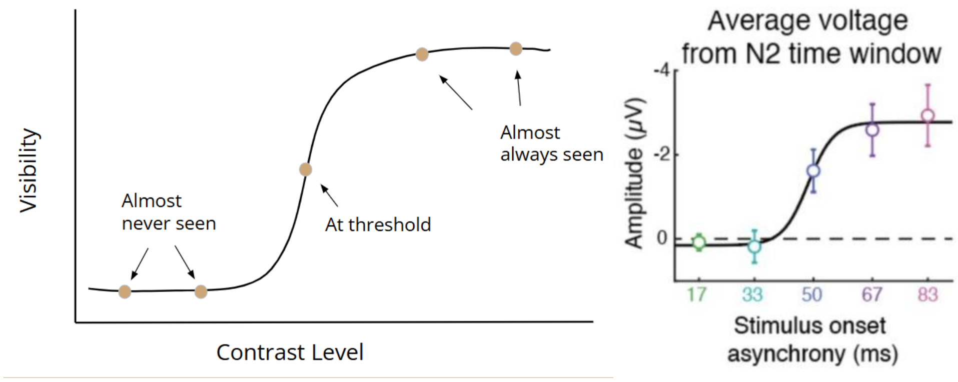

Figure 2. Examples of a bifurcated curve.
A diagram of a bifurcated S-curve (left) alongside the same pattern mapped for the N2 wave as amplitude (strength of signal response in participants) across stimulus onset asynchrony (visibility of the stimulus) (Cohen et. al., 2024).

1.5.2. Future Directions and the Purpose of this Thesis

There was one complication with the study design which prevented precise isolation and quantification of the N2 signal–the use of backwards masking. The EEG data indicated additional ERPs evoked not just by the stimulus but by the ensuing masks themselves, which was sought to be accounted for by implementing mask-only trials that did not present a stimulus beforehand, as in (Del Cul et al., 2007). The idea was to isolate which ERPs were evoked by the masks and isolate them from the ERP data of the stimulus+mask trials, though unfortunately this subtraction process was imperfect. As a result, while the N2 was measurable, the signals between it and the P1 were essentially uninterpretable and leftover artifacts from the mask-evocations were still visible. So, while the procedure was sufficient to indicate the presence of a perceptual awareness ERP distinct from neighboring activity events, it yielded what was at times an obscured and imprecise measure of activity across the full relevant timeframe. It also remains unclear whether the VAN, observed in the previous study, bifurcates alongside the fcN2 or if it is more continuous and linear.

This thesis continues the EEG-based multifactor task paradigm in pursuit of isolating the N2 wave as a novel perceptual awareness NCC. The data collection, completed in tandem by myself and Angelica Nicolacoudis, utilizes the no-report condition and non-linear contrast detection threshold methodology used by the previous study while simplifying the stimulus presentation such that content-specific activity will not overcomplicate the data. After data collection, I aggregated and analyzed the data for ERPs, focusing on time windows and electrode pools which have emerged bottom-up as points of interest in the data. The overall initial aim of this research was to isolate neural events–particularly the N2–as a temporal, event-related, neurologically independent event necessary for the generation of awareness as an experience of consciousness. An additional aim was to further develop an understanding of whether the VAN corresponds more strongly to sensory processing of the stimulus or to conscious awareness.

The significance of the N2 in particular lies in its temporal position–it occurs markedly after the window of local sensory processing (Koch et al., 2016; Lamme, 2018), but as a post-perceptual processing awareness signal it occurs earlier than a prefrontal-exclusive theory of higher-order thought processing would predict (Dehaene et al., 2017; Odegaard et al., 2017). This position suggests that the wave could offer insight into the signaling relationship between lower- and higher-order cortices during initial feedforward and subsequent top-down communication, offering a link in the timeline of recurrent processing that would lend support to models which posit this transmission to a global network. This could potentially further illuminate at what point in this process conscious perception is activated and “accessible.”

PART 2: METHODOLOGY

2.1. Experimental Paradigm

2.1.1. Stimulus of Choice: Gabor gratings

The experimental approach differs most significantly from the previous paradigm in the use of a simplified visual stimulus. Instead of faces and masks, participants were presented with “Gabor” gratings–circular images which contain diagonal sine-wave lines facing either 45 or 135 degrees.

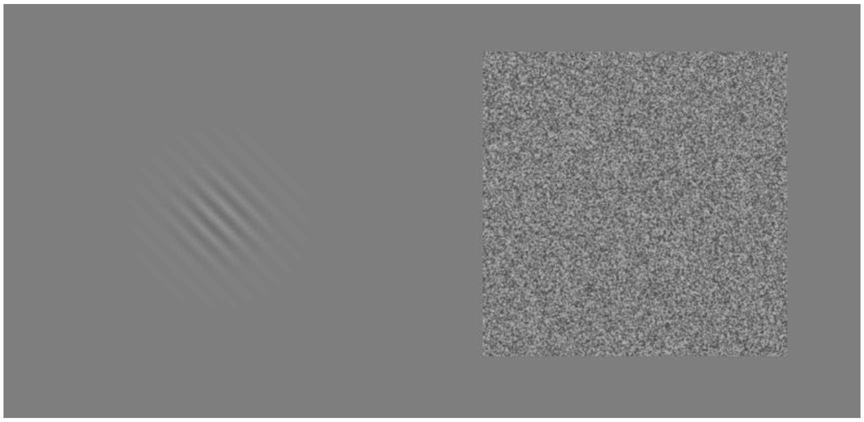

Figure 3. Gabor gratings.

The Gabor grating stimuli used in the procedure (left) and a static image of the visual noise in which they were embedded (right).

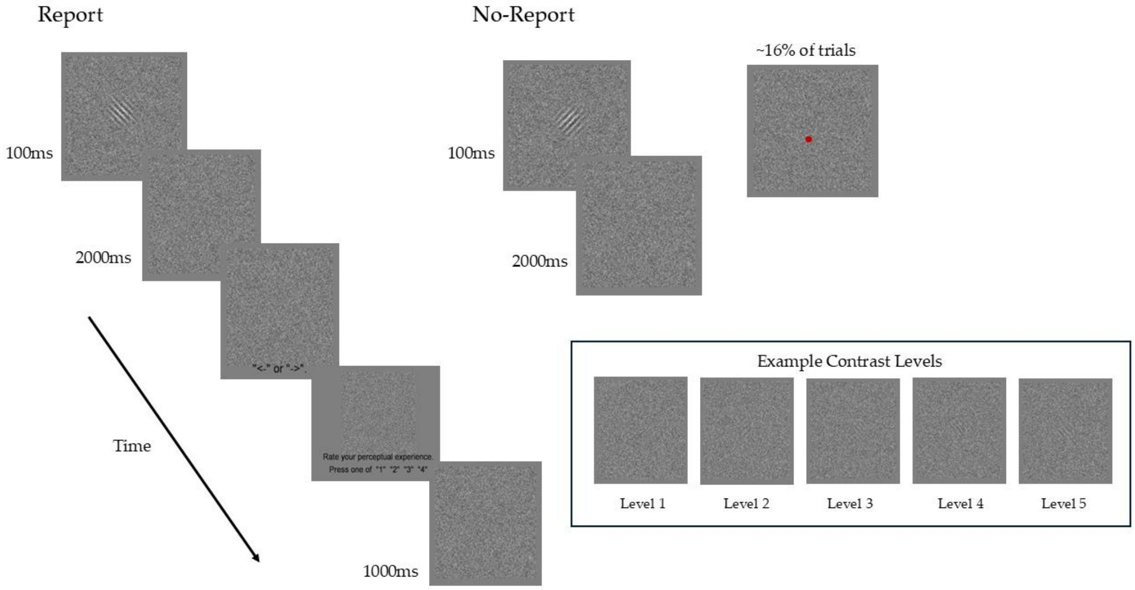

Figure 4. Diagram of the stimulus presentation.

A visual demonstration of how the Gabor stimuli were presented over time, with example Gabors at 5 different linearly spaced contrast levels. In the report (left), subjects were asked two prompts after being presented with the stimulus. In the no-report (right), their task was to detect a red dot probe unrelated to the Gabors. Objective contrast value of the Gabors was variable depending on the threshold modifier calculated for each individual subject.

These gratings were embedded in continuous "dynamic noise" (visually analogous to TV static) which was used to obscure the gratings. These stimuli are singular and simpler in content than using faces quickly obscured by subsequent masking. The purpose of this simplification was to only produce one event-related potential for each presentation of the stimulus, leading to more consistent and non-obfuscating signals that can more easily be distinguished from neighboring events.

2.1.2. Contrast thresholding.

Each Gabor was presented for 100ms, so the duration was not manipulated–instead, the contrast level was adjusted around an established threshold. The first segment of the experiment was a 3-block *quest procedure*, in each of which the participant was presented with 50 gratings. For each presentation of the grating, the participant was tasked with indicating whether it faced upward to the right or left using corresponding arrow keys. The quest segment operated via a staircase procedure which adjusted the contrast/visibility depending on the accuracy of the participant. If they were consistently guessing correct, this indicated that they could see the gratings most of the time, and the contrast was decreased. When they guessed incorrectly, the contrast would increase until they could guess above chance. Over 150 total trials, the gratings were calibrated to establish a *visual threshold*, at which subject could guess the orientation with ~75% accuracy. In this case, 75% accurary indicates 50% awareness, as they had a 50% chance of guessing correctly even in the half of trials where did not see the stimulus (50% accuracy In the 50% of trials which are unseen = an additional 25% accuracy overall).

In the effort to achieve 75% average accuracy at threshold using the quest procedure, pilot testing revealed that there was a learning effect which results in higher real accuracy among participants than what the quest attempts to achieve. As such, various measures were implemented over the course of pilot testing to correct for this effect. First, we excluded the first block of 50 trials from the final accuracy measure to account for the over-time improvement in detection skill. This improved the average accuracy in the real experiment somewhat, but not up to 75%. Further measures were taken by adjusting the procedure such that it would only calibrate to 70% accuracy, and finally only 65%. Both decreases further increased the efficacy of the calibration for the actual test. Once this threshold value was established, the Gabors for the remaining procedure were spaced at five linear contrast values, indexed by arbitrary integers, such that they were rarely or never visible for the bottom two and almost or always visible for the top two. In addition to these contrasts, there were periodically blank trials which did not present a grating. After the first five pilot tests, a sixth contrast level, higher than the fifth, was added to ensure that a point of bifurcation could still be seen even if the participant’s threshold level was higher during the actual procedure than what was calculated during the quest. This sixth contrast level was ultimately removed for the final experimental design, as data from the complete pilot set revealed bifurcated signal patterns around threshold. Therefore, for the actual experiment there were a total of 6 stimulus types (five contrast levels and a blank).

2.1.3. Report and no-report.

After establishing the participant's visual threshold, they were then tasked to complete the two main segments of the experiment: report and no-report conditions. Initially, both segments consisted of 18 blocks of 42 trials (126 trials per contrast level), both completed during the same experimental session under EEG measurement. However, with the addition of the sixth contrast level, the experiment reached a length at which the participant needed to actively engage for >2 hours; this led to negative feedback and concerns that the data would be compromised by fatigue. For the next several participants the experiment was split across 2 sessions on two different days.

Report Condition

Finally, we determined that sufficient pilot data from the report segment would serve as a suitable measure, and that this segment could be ran as a behavioral-only segment without EEG in the final experimental design. This justified reducing the length of the condition significantly, from 126 to just 40 trials per condition. This reduced the length of the condition from 45+ to ~15 minutes, allowing the whole experiment to be ran in a single session.

In the report condition, 600ms after the onset of each stimulus, the participant was prompted to report whether the grating faced left or right (the same task as the quest procedure). Subsequently, they were asked to give a perceptual judgment rating of the clarity of the stimulus on a 4 point scale:

1: No stimulus visible;

2: Faint awareness that the stimulus was presented;

3: Moderate awareness of stimulus and direction;

4: Confident awareness of stimulus and direction.

The PAS (Perceptual Assessment Scale) allowed for a subjective measure of perceptual awareness, which could be compared to the objective forced-decision measure in the left-right report task. In this system, a trial was considered "visible" If the participant gave a score of 2 or above on the scale. The intention was that this value would be 50% at the visual threshold value (contrast level 3 of 5).

No-report condition

In the final experimental design, EEG recordings were only taken during the no-report condition, which composed the bulk of the session after the report segment was reduced. In this condition, the stimuli were presented without interruption, each appearing sometime between 800 and 1200ms after the onset of the previous. Participants were instructed not to actively pay attention to the Gabors or take mental note of them, so as to minimize the activation of the P3 signal. To maintain participant attention at the general area of the gratings, they were given a task related to an additional visual stimulus. This was initially a "coherent motion probe" in the dynamic noise, in which the static would suddenly shift in a synchronized motion upward. This led to somewhat unclear data in the pilot trials, which was attributed to the sensory similarity between the motion of the static and the Gabor grating appearances.

The probe stimulus was later changed to brief red dots, the rationale being that small colored dots would be more visually distinct than another kind of greyscale motion. Per block of trials, 6 red dot probes would appear, each within the same window around the center of the noise where the gabors would appear.

2.2. Logistics and Parameters

2.2.1. Pre-registration.

Planned methodology for procedural data collection and analysis was pre-registered on the Open Science Framework (OSF).

2.2.2. Exclusion.

Various anomalies can exclude participants from data analysis, such as an unusually high false alarm rate (stating a gabor was seen during a blank trial). The idea of noting false alarm rate is to determine if the participants' responses during stimulus presentations are meaningfully distinct from how they respond to pure noise, informing the degree of confidence with which their results can be interpreted. Additionally, while a 75% accuracy threshold is ideal, previous studies on bifurcation dynamics have allowed variability in a range of 65-85% accuracy and seen a conserved non-linear curve nonetheless (Sergent et al., 2021). It is furthermore possible that at the lowest contrast levels the participants will guess above chance even if they claim not to consciously notice the stimulus in their report ratings. Ultimately, it was decided that participants would be excluded if they failed to meet any of the following criteria:

1. 5%+ objective increase in 2AFC accuracy from contrast level 2 to 3
2. 5%+ objective increase in percent-seen reported by the PAS scale from contrast level 2 to 3
3. 5%+ objective increase in 2AFC accuracy from contrast level 3 to 4
4. 5%+ objective increase in percent-seen reported by the PAS scale from contrast level 3 to 4
5. A false alarm rate which does not exceed 30%
6. At least 65 trials remaining in each condition after artifact rejection and corrective measures during pre-processing

2.2.3. Participant Collection.

30 participants worth of usable EEG data was gathered from individuals between the ages of 18 and 24, sourcing primarily from the Reed College undergraduate student body. 43 total participants were collected and 13 were excluded (1 for false alarm rate, 6 for low 2-3 difference, 6 for low 3-4 difference). Initial pilot data consisted of 10 subjects plus 2 exclusions, collected separately from the designated experimental trials. Experimental procedures were adjusted significantly after the first five pilots for the no-report section, but all ten completed an EEG report section which was viable for analysis. In total, 55 people participated In the experiment.

2.2.4. Approach Timeline.

Initial efforts were focused on troubleshooting the piloted code and procedure over the first four weeks. Time-testing ensured accurate measurement and recording of the hard- and soft-coded procedure elements in Matlab. Behavioral piloting was performed to ensure the participant experience was as-intended and the quest thresholding procedure was performing its operative function as intended. EEG piloting then ensued for the remainder of the semester, troubleshooting as necessary in tandem with the collection of pilot data. Four participants were gathered in December as potential “true” subjects following pilot analysis. In January, final exclusion criteria were established, along with the final procedural adjustment of opting to forgo the sixth contrast level. Participant testing ensued until April 1st, 2025, when the thirtieth useable subject was collected, and the dataset was completed.

2.3. Analysis Strategy

2.3.1. Behavioral Analysis.

Collaborative analysis was conducted concurrently with the collection of pilot participants to determine if the participants are successfully producing the non-linear bifurcated patterns of awareness anticipated in the accuracy and percent-seen measures of the report trials. The 2-alternative forced choice measure and 4-point PAS scale in the report condition has allowed for continual procedural adjustment with the goal of aligning the contrast levels such that they are as close as possible to 50% seen (and guessed with 75% accuracy) at threshold, with the two shortest intervals nearly never seen and the two longest intervals almost always visible. While EEG measurements needn’t be collected during the report segment, the behavioral measures taken from this condition were vital for assessing the viability of the no-report measurements and thereby validating each participant’s data.

2.3.2. Event-Related Potentials.

Following data collection, the EEG data collected from the no-report condition was examined across averages and among the individual participants for the presence of activity signatures at particular times. The previous paper gave a reference point for anticipated ERPs as well as areas of interest. An initial positivity (P1) was nticipated from 100-140ms during the sensory processing phase, and the P3b was measured from 300ms and beyond. They measured an early VAN, occurring from 140-200ms at electrodes 54 and 57. With the collection of pilot data and a gradually building grand-average of experimental subjects, there were adjustments and additions to these initial time windows, though the fcN2 selection area was unchanged. The final zones of interest (Figure 4) were as follows:

VAN: 250-350ms, electrodes 42-44, 54-56, 27 & 28 (which corresponds to PO3, POz, PO4, O1, Oz, O2, O9, Iz, O10 of the standard 10-20 system)

fcN2: 250-300ms, electrodes 1-11, 17-19 (which corresponds to CPz, CP1, CP2, Cz, C1, C2, C3, C4, FCz, FC3, FC4, Fz, F1, F2 of the standard 10-20 system)

Late Positivity: 450-650ms, electrodes 14-16, 28-30 (which corresponds to CP3, P1, Pz, CP5, P5, P3, PO3 of the standard 10-20 system)

Late Negativity: 450-650ms, electrodes 20&21, 34-37, 49, 50 (which corresponds to AF7, Fp1, Fpz, Fp2, AF8, AF3, AFz, AF4 of the standard 10-20 system)

P3b: 300-500ms, electrodes 1, 4-6, 13-15 (which corresponds to Pz, P1, P2, CPz, CP1, CP2, Cz of the standard 10-20 system)

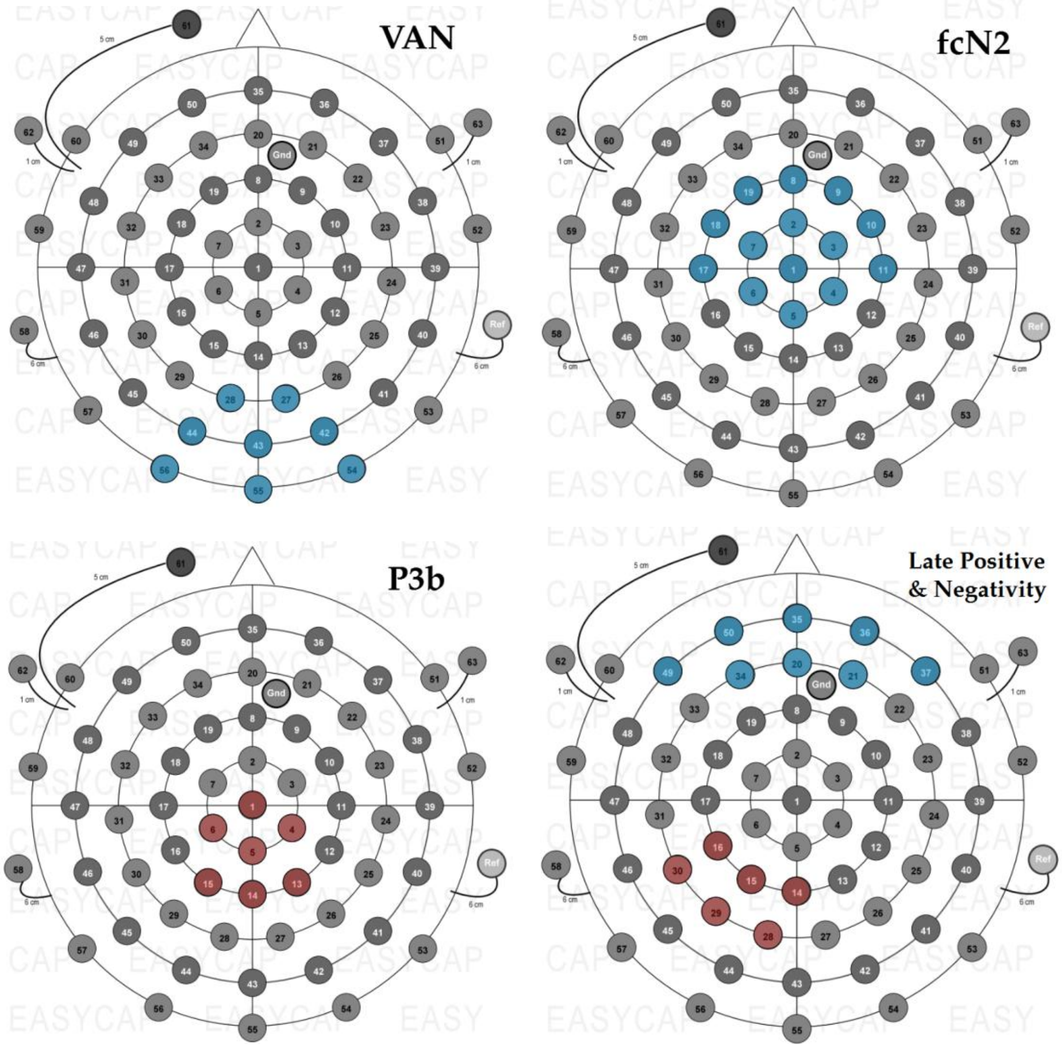
VAN

fcN2

P3b

Late Positive & Negativity

Figure 5. Diagram of electrode pools selected for each ERP.
ERPs are ordered left-right and top-bottom by temporal order of appearance.
Negative signals are blue and positive signals are red.

The Visual Awareness Negativity in the posterior cortex, expected to peak at ~200-250ms, was of significant interest in this experiment. Observation of the VAN would offer insight into its trajectory and linearity as it corresponds to stimulus strength. As the dataset of participants grew, time windows and electrodes of interest were continually discovered and assessed accordingly. Of interest for each ERP was whether it followed a bifurcated or linear trajectory (Kouider et al., 2013), indicating whether it correlates directly to sensory strength as opposed to conscious awareness.

EEG Data Processing

All EEG raw data was exported to Brainvision Analyzer for processing. Stimulus trigger codes were time-correcte and a high-pass filter of 0.1Hz was ran over each file. The bipolar vertical and horizontal electrooculogram channels (VEOG and HEOG) were utilized to semiautomatically detect blinks and eye movements, which were eliminated alongside other artifacts using a combination of Independent Component Analysis (ICA), reverse-ICA, and semiautomatic artifact correction/rejection. This process order allowed for corrective measures which preserved as many trials as possible without conserving their artifacts. Baseline correction was used to normalize the amplitues across trials and prevent excessive fluctuation-induced variability. After pre-processing, the data was segmented by condition (contrast level). The data was grand-averaged as the subject pool grew using the software's ability to aggregate mean data values and visualize them. The data was also separated into pools of electrodes depending on which ERP was being analyzed. All graphs and plots were made using Matlab.

PART 3: RESULTS

3.1. Behavioral Results

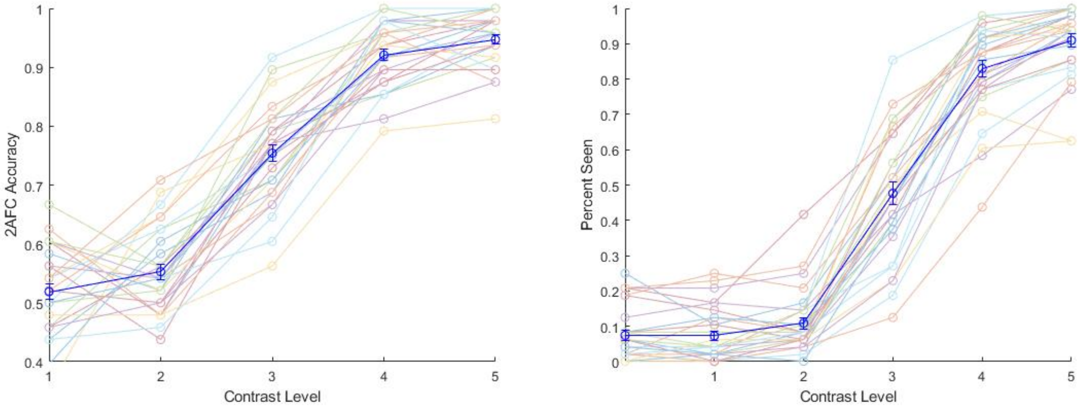

Figure 6. Behavioral response patterns from the report condition.
Average of participant (n=30) responses for both report condition tasks at each contrast level. Faint lines are individual subject data from which the averages were yielded. Objective accuracy in the 2AFC (Left) and subjective awareness (right) both show significant differences between contrast levels 2 and 3 as well as between 3 and 4. Subjective awareness was a measured as a binary product of the PAS scale; it is the percent of responses in the 4-point scale which were 2, 3, or 4 vs. 1 (i.e., something seen vs. nothing seen).

Behavioral results were used as exclusion criteria for participants throughout. As such, both curves follow accuracy/rating patterns which emulate the intended sigmoid effect. These are indications of overall experimental validity, as the ERPs in the EEG data can be compared to a successfully designed threshold pattern of visibility. The 2-alternative forced choice left/right discrimination test (Figure 5, left) after each Gabor yielded noticeable jumps in the steps out from contrast level 3 toward 2 and 4, but not between 1 and 2 or 4 and 5. The same effect was produced in the four-point subjective visibility rating values when plotting the percent-seen as a product of the percent of response values which exceeded 1 (as "1" indicates no stimulus seen while 2-4 indicate some degree of visibility) (Figure 5, right).

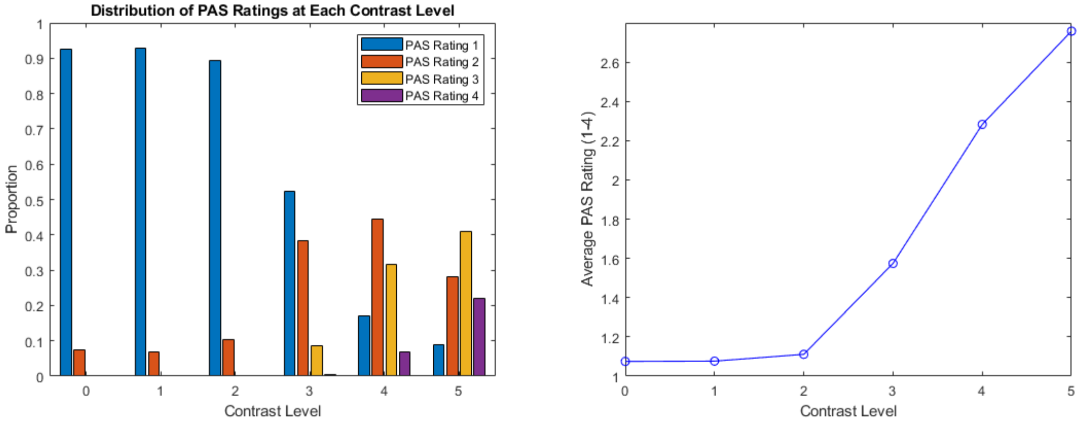

Figure 7. Distributions of the subjective visibility assessment.
A (left) indicates, at each contrast level, what proportion of total trials each rating value composed. B (right) maps the average subjective rating given by participants for each contrast level.

For blank trials as well as contrast levels 1 and 2, participants almost universally gave a rating of "1," indicating they did not see any stimulus presented. The proportion of "2" ratings did not change significantly between the bottom two contrast levels and the blanks, indicating no difference between the response elicited on average by these signals versus pure noise (Figure 6, left). A mix of ratings appear at contrast level 3, with participants marking above a "1" almost exactly 50% of the time. At contrast level 4, ratings of 2 and 3 were both more frequent than "1"/unseen, and at the highest contrast level participants were less likely to give a 1 than any other rating.

Average PAS rating, plotted by contrast level, does not change from contrast level 1 to 2 but increases linearly from threshold and onward (Figure 6, right). This is a change from the bifurcated dynamic that appears in the objective task data (2AFC orientation discrimination) and in the subjective rating data when a binary split is made between unseen (PAS = 1) and seen (PAS = 2, 3, or 4). This difference reflects the nature of the perceptual assessment scale itself, i.e., it was designed to capture more graded changes to perceived clarity. When the participants did not see the Gabor, regardless of its contrast, they pressed 1 to indicate no stimulus detected. Once the stimulus was high enough in contrast for them to see it, their average PAS ratings increased for each contrast level. Both observations are indicators that the PAS scale was, on average, performed as intended by the participants, indicating that (1) their ratings reflected their perceptual experiences and (2) their perceptual experiences were as intended.

3.2. Event Related Potentials

3.2.1. The P1 and P3b

There was no P1 measured in the final grand-average data across the 30 subjects (Figure 7, left). There was also essentially no posterior positivity in the no-report condition during the suspected P3 time window (Figure 7, right), which is another indicator of experimental validity. A P3 would have likely been the strongest of the signals by far and could have drowned out the other temporally overlapping or neighboring activity events. The P3 here is measured in an adjusted time window (300-450ms) to account for the emergence of a new positive posterior signal which displays notably different qualitative pattern topography and signal behavior.

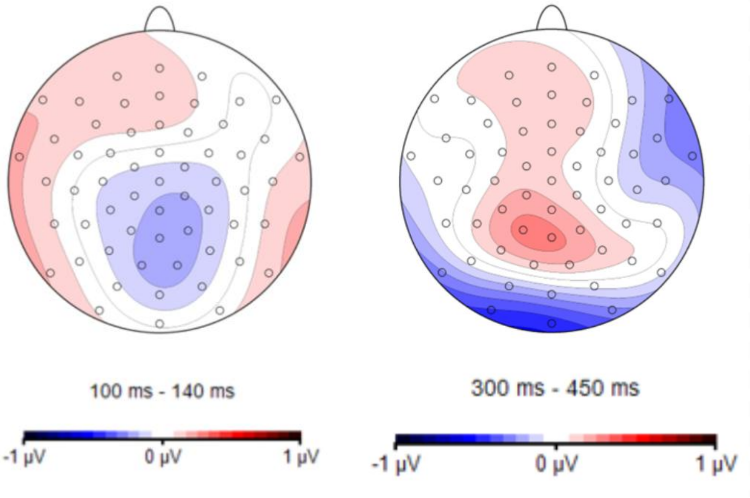

Figure 8. Scalp map of the P1 and P3 time windows.
Diagram of scalp topography spanning from 100-140ms and 300-450ms. Both maps have had their scales manually set from -1 to 1 microvolt.

3.2.2. The VAN

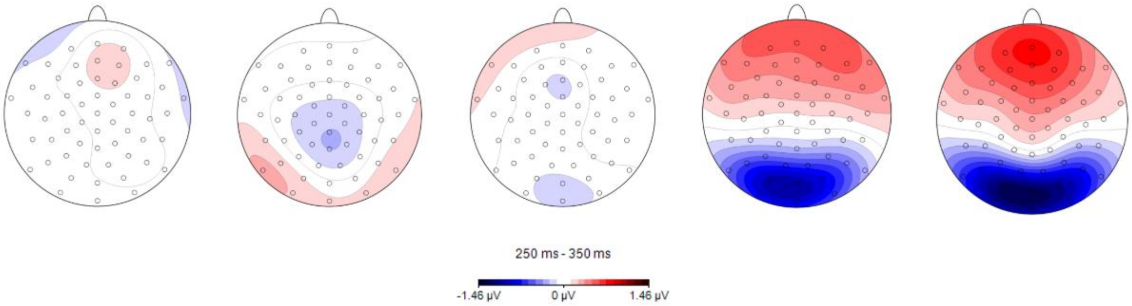

Figure 9. VAN scalp maps at each contrast level.
The posterior negativity of the VAN first appears at threshold and is significantly stronger at levels 4 and 5.

Based on the scalp topography of the signal across contrast levels (Figure 8), it seems the VAN was almost exclusively elicited by the fourth and fifth contrast levels, though it does show an early trace at level 3. Statistical analysis revealed that while there was no significant difference between contrast levels 2 and 3 (*t*(29) = 0.87, *p* = .39), there was a significant increase from 3 to 4 (*t*(29) = 4.19; *p* < .001). There was also no significant difference between levels 1 and 2 (*t*(29) = -0.43; *p* = .68) or between 4 and 5 (*t*(29) = 1.28; *p* = .25). Observation of both the waveforms and the mean amplitude values of the electrode pool indicates significantly greater activity elicited by the two highest contrast levels and almost negligible activity from the lower three (Figure 9).

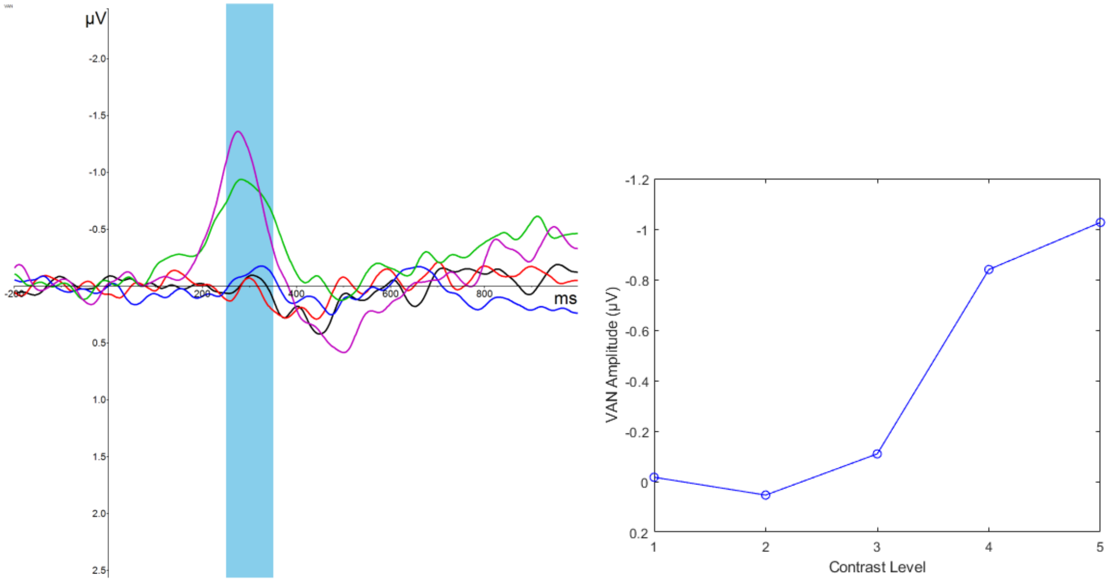

Figure 10. Amplitude patterns of the Visual Awareness Negativity.
Left: the waveform map of the VAN electrode pool at each contrast level (black = 1, red = 2, blue = 3, green = 4, purple = 5). Measured area is highlighted in blue.
Right: the mean amplitude values of the VAN electrode pool plotted across contrast levels. Only the jump from 3-4 was significantly different (*p*<0.001).

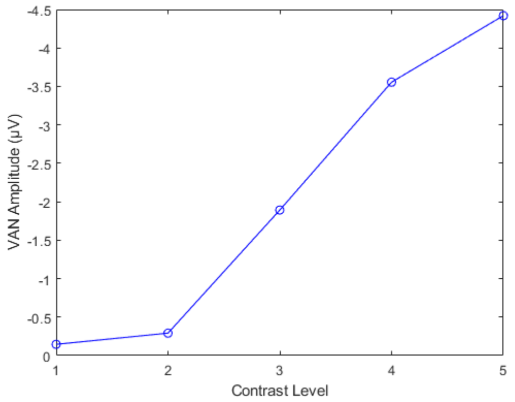

Figure 11. Amplitude patterns of the VAN In the pilot report condition. Mean amplitude values of the pilot subjects (n=10), who were the only subjects to receive EEG recordings in an extended report condition.

While the shortened report segment in the main experiment was used only to establish a behavioral point of reference, the pilot subjects did undergo EEG recordings during their extended version of the condition. From this pool of 10 pilots, a VAN was visible from 250-350ms in the electrode pool ultimately used for the main experimental analysis. Mean amplitude values of this pool plotted across contrast levels (Figure 10) demonstrate bifurcation, with significant jumps from level 2 to 3 ($t(9) = 4.14$; $p = .00251$) and 3 to 4 ($t(9) = 5.88$; $p < .001$) but not from 1 to 2 ($t(9) = 0.78$; $p = .45$) or 4 to 5 ($t(9) = 1.97$; $p = 0.08$).

3.2.3. The Fronto-Central N2

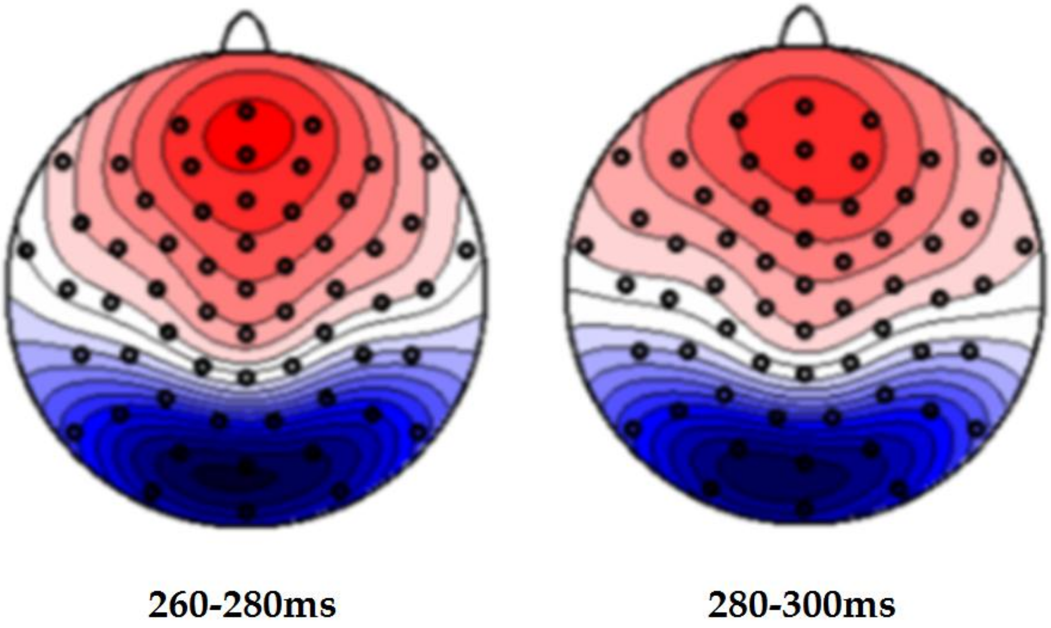

Figure 12. Grand average scalp maps in the fcN2 time window. There is no recognizable negativity in the fronto-central region of the scalp during the time window in which this ERP has previously been visible.

One of the most intriguing developments of the experiment was the complete lack of any fronto-central negativity in or around the anticipated time window (Figure 11). While the VAN was highly visible at this time, it became clear from the pilot data as well as the first 12 subjects that a negativity in the central electrodes (1-11, 17-19) would not be evoked by our experimental conditions. Subsequent research into the nature of the fcN2 involved a literature investigation pursuing the variables manipulated across no-report EEG studies which did and did not evoke the signal.

3.2.4. Late Signals

An intriguing finding was the appearance of two signals in a late 450-650ms time window. One is a very frontal negative signal which will be labeled the Late Frontal Negativity (LFN) (Figure 12) in this report, while the other is a posterior, somewhat lateralized positive signal which we label the Late Visual Awareness Positivity (LVAP) (Figure 13).

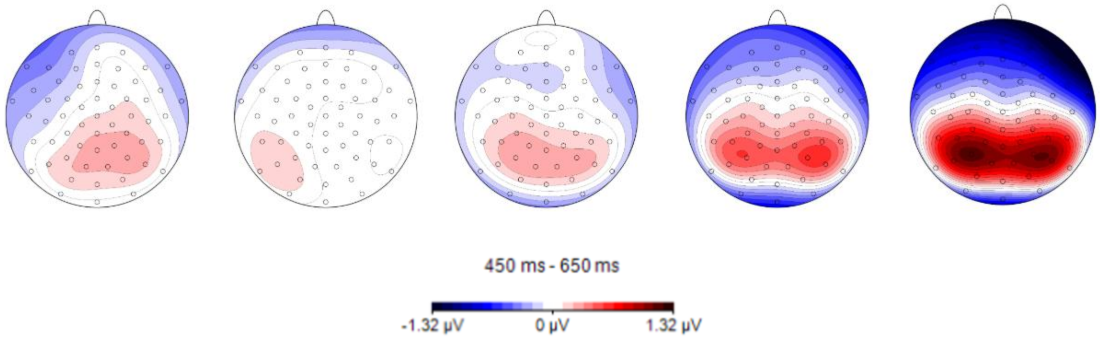

Figure 13. Late scalp topography at each contrast level.
Scalp maps of the late time window (450-650ms) at each contrast level (1-5, left to right). Scale automatically plotted between -1.32 and 1.32 microvolts.

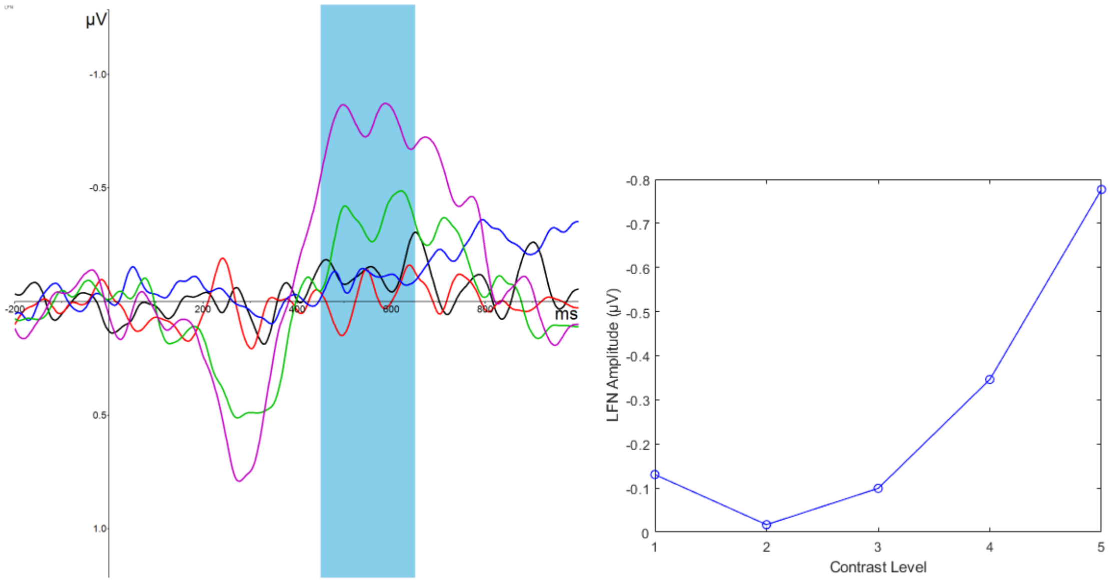

Figure 14. Amplitude patterns of the Late Negativity.
Left: Grand average waveforms of the selected electrode pool at each contrast level (black = 1, red = 2, blue = 3, green = 4, purple = 5), demonstrating a negative spike around 450 ms. The area measured is highlighted in blue. Right: mean amplitudes plotted across contrast levels.

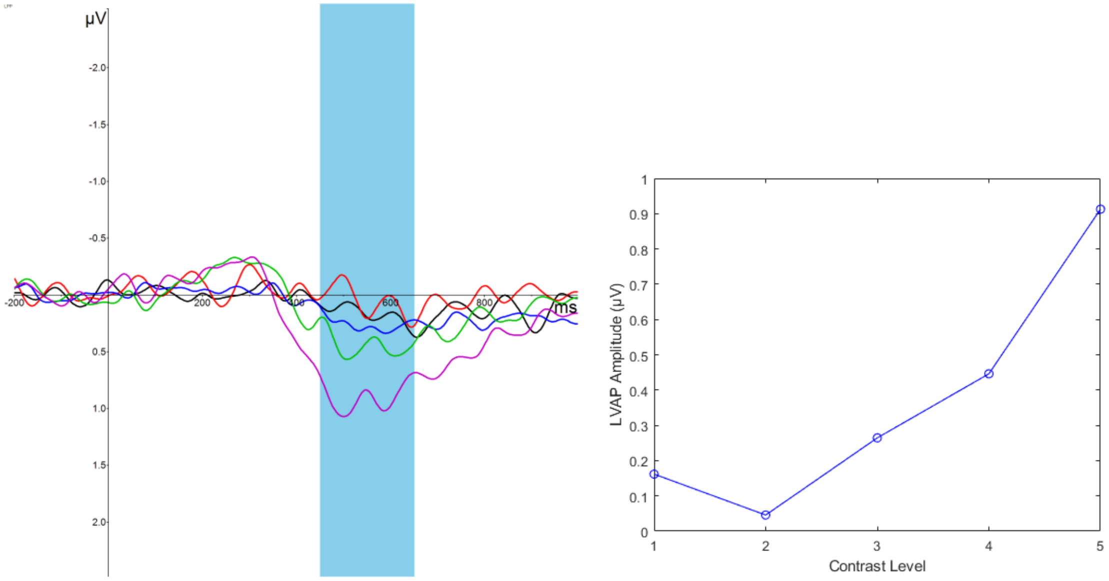

Figure 15: Amplitude patterns of the Late Positivity.

Left: Grand average waveforms of the selected electrode pool each contrast level (black = 1, red = 2, blue = 3, green = 4, purple = 5), demonstrating a positive spike around 450 ms. The area measured is highlighted in blue. Right: mean amplitudes plotted across contrast levels.

The two signals demonstrate similar patterns of mean amplitude value increases. For both, contrast levels 1 and 2 did not induce significantly different levels of activity (LVAP: *t*(29) = 0.90; *p* = .38; LFN: *t*(29 = -0.66; *p* = .51), nor did 2 and 3 (LVAP: *t*(29) = -1.88; *p* = .07; LFN: *t*(29) = 0.59; *p* = .56) For the LVAP, level 4 induced significantly greater amplitude activity than level 3 (*t*(29) = -2.10; *p* = 0.046748524), while for the LFN it did not (*t*(29) = 1.50; *p* = .14). For both, level 5 was significantly stronger than level 4 (LVAP: *t*(29) = -3.82; *p* < .001; LFN: *t*(29) = 2.37; *p* = 0.0246). Both signals are of relatively similar strength at each contrast level and appear to demonstrate an exponential growth pattern as contrast linearly increases. Both signals also present two peaks each within the measured time window, the first at ~495ms and the second at ~590ms.

Dissection of Bilateral Positivity Points

The LVAP appears to originate at two different points, one on each side of the posterior scalp. This prompted subsequent individual observation of each subject's scalp topography, which revealed multiple variations of the signal across participants.

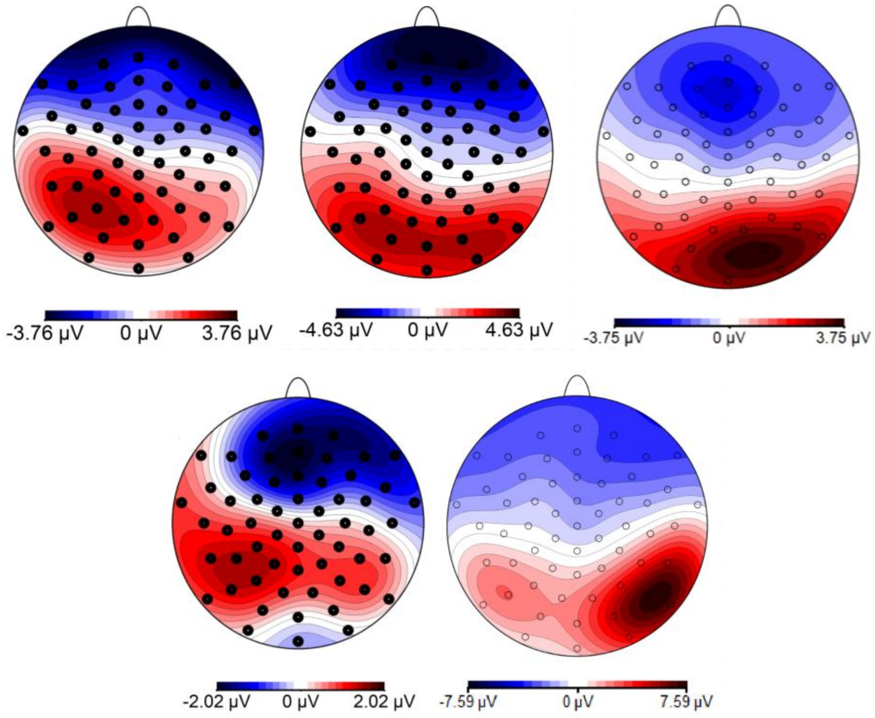

Figure 16. Individual Topographical variations of the LVAP.
Individual subjects demonstrating variation of source point(s) of the positivity originating in the posterior half of the scalp at the 450-650ms. From left to right, then top to bottom: left only, both (even strength), right only, stronger left, stronger right.

Across the thirty subjects, there were four for whom there was no positivity in the candidate time window which originated from the posterior half of the scalp. Only those whose recordings revealed this signal were considered to have an LVAP. 21 subjects presented a left-originating positivity, while 18 presented a right-originating one. 13 subjects produced both, giving a 50% bilateral positivity rate among those who present the signal.

The lateral patterning warranted examination of each positivity source individually. By determining the strongest sources of positivity in each subject with an instance of a posterior positivity on one side or the other in the 450-650ms time winow, new groups of electrodes were pooled together for the left and right manifestations of the LVAP (Figure 12).

Left: 44-46, 28-30 (corresponding to O1, PO7, P7, TP7, PO3, P5, CP5)
Right: 40 & 41, 25 & 26, 53 (corresponding to P6, P8, P10, CP6, TP8)

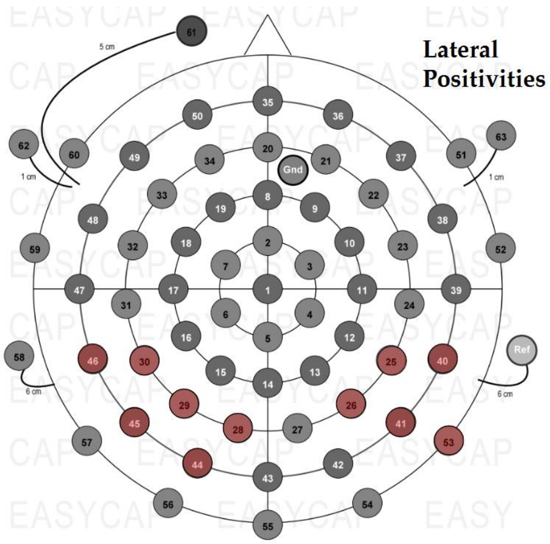

Figure 17. Left and right posterior positivities.
Illustration of the electrodes selected on the left and right sides based on measurement of subject-by-subject lateral positivity signal source points in the 450-650ms range.

In isolation, the left and right positivities both maintain an exponential rate of mean amplitude increase, with the left in particular approximately doubling in amplitude with each contrast increase. Separating this ERP by lateral point of origin does not appear to create any significant change in its signal pattern around visibility threshold.

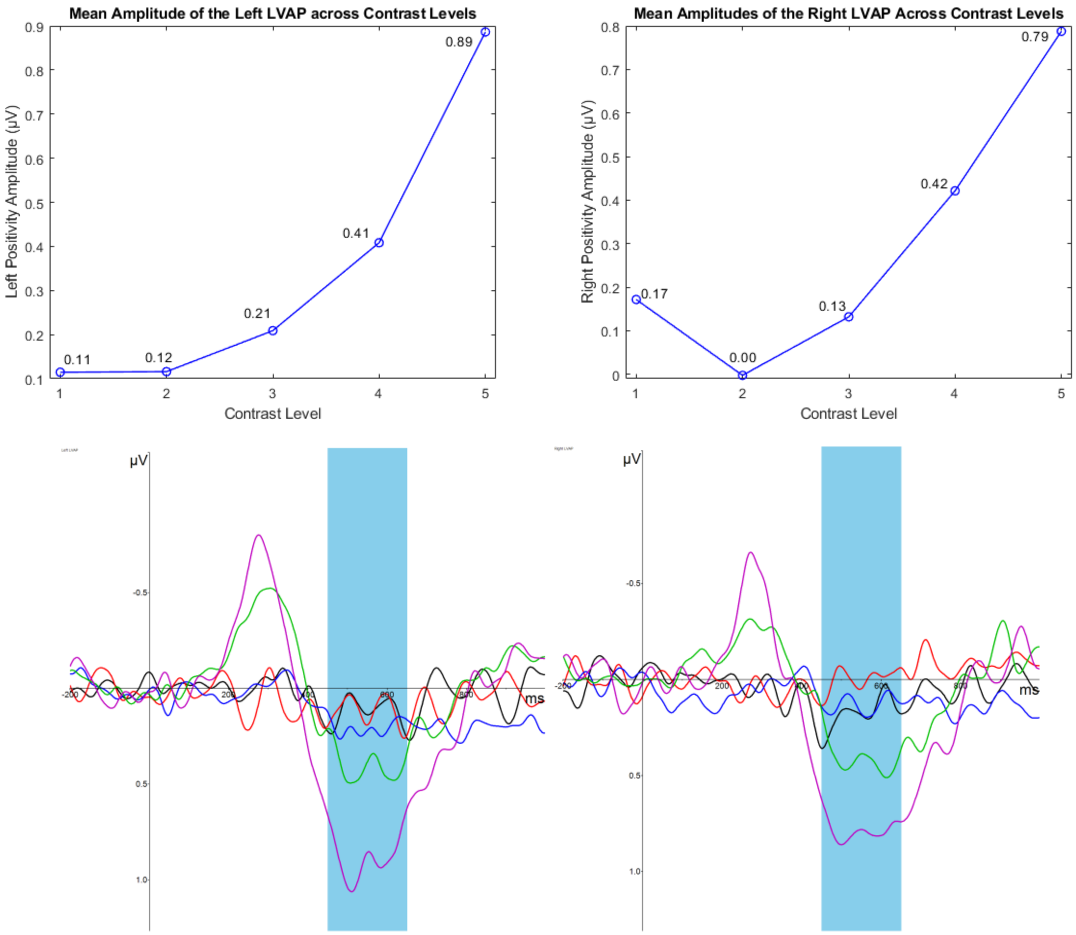

Figure 18. Amplitude patterns of the Left and Right Late Positivities.
Above: Mean amplitudes plotted across the contrast levels for the left (left) and right (right) positivity source points in the posterior scalp region throughout the 450-650ms time window. Below: Grand average waveforms of the left and right LVAP electrode pools at each contrast level (black = 1, red = 2, blue = 3, green = 4, purple = 5). The area measured is highlighted in blue.

These left and right partions demonstrate similar patterning to the combined positivity, and also present two peaks within the measured window. The right LVAP also shows peaks at ~495 and ~590ms, while the left positivity peaks in a slightly tightened time window at ~510 and ~585ms.

PART 4: DISCUSSION

4.1. Multiple Interpretations of the VAN

4.1.1. Pattern Dynamics

One of the principle aims of this research was to achieve a better understanding of the Visual Awareness Negativity ERP by determining whether it presents a bifurcated pattern dynamic of signal strength around visual threshold. Ultimately, the signal yielded an elusive variation of an ordinary sigmoid curve which lends itself to a variety of possible explanations.

A bifurcated curve would expect to see the following statistical pattern:

- No significant difference between contrast levels 1 and 2
- Significant increase from level 2 to 3
- Significant increase from 3 to 4
- No significant difference between 4 and 5

The strength of the VAN did not significantly vary between contrast levels 1 and 2, nor between levels 2 and 3. There was, however, a significant increase from level 3 to 4. Finally, there was no variation between levels 4 and 5.

4.1.2. The Effect of No-Report on Calibrated Visibility

There are at least two ways this data can be interpreted. From a statistical perspective, it appears that the VAN *does* bifurcate (as in, there is a singular jump between two unchanging amplitude ranges), but its threshold level is not at contrast level 3. One way to reconcile this would be to concede a methodological limitation of the no-report paradigm.

The amplitude of the VAN was ~3-4x weaker in the no-report than the report condition. One possible explanation of this difference is that participants were simply not paying as much attention to the stimuli in the no-report condition (after all, they were instructed to be on the lookout for occasional red dot targets, which may have drawn too much attention away from the gabor gratings). Given this, it follows that stimuli of the same contrast strength might be seen less often in the no-report condition than in the report condition, including the quest procedure where the contrast values are calibrated (which also required trial-by-trial reports of stimulus orientation). This change in the contrast threshold due to differential attention is supported by psychophysics literature which has demonstrated that spatial attention does have an effect on the perceived contrast of a stimulus (Carrasco, 2018). Perhaps the stimuli need to be of a higher contrast in order to reach threshold in a no-report condition with less task-based attention devoted to the stimuli, leading to a shift in which a “true threshold" is reached somewhere between the third and fourth contrast levels (Figure 12).

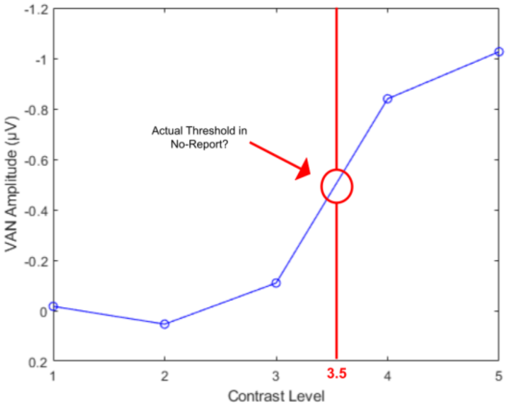

Figure 19. A threshold-shifted interpretation of the VAN.
Conceptual illustration of what may be indicated by statistical analysis of VAN mean amplitudes.

Though there were only 10 subjects in the report condition, the potential limitations of the no-report condition warrants returning to this pilot data for insights. In this condition, there were significant increases from both levels 2 to 3 and 3 to 4, but not between any other levels. Statistically, this matches the expected outcome of a bifurcated sigmoid dynamic. This lends further credibility to the idea that the VAN is indeed a neural correlate of conscious awareness, and that its statistical variations in the no-report condition are simply reflective of how this method affects awareness. In any case, the NCC candidate status of the VAN is strengthened by these findings.

4.1.3. Reconsidering the Bifurcated Model of Awareness Correlates

Based purely on observation of the amplitude plots in both conditions, there is an additional possible characterization of the VAN's signal behavior. In both, there was less statistical confidence in the lack of amplitude increase from level 4 to 5, where the signal was almost certainly seen, as opposed to levels 1 and 2 where it was almost certainly not. If the report condition is indeed a more accurate indication of the VAN's behavior, this observation is even more supported as the change from 4 to 5 was extremely close to achieving statistical significance ($P = 0.08$). It can't be said whether this effect would be enhanced or mitigated with a greater sample size, but as it stands, the curve bears an intriguing visual similarity to one other metric plotted across the contrast levels: the average PAS scale ratings (Figure 13).

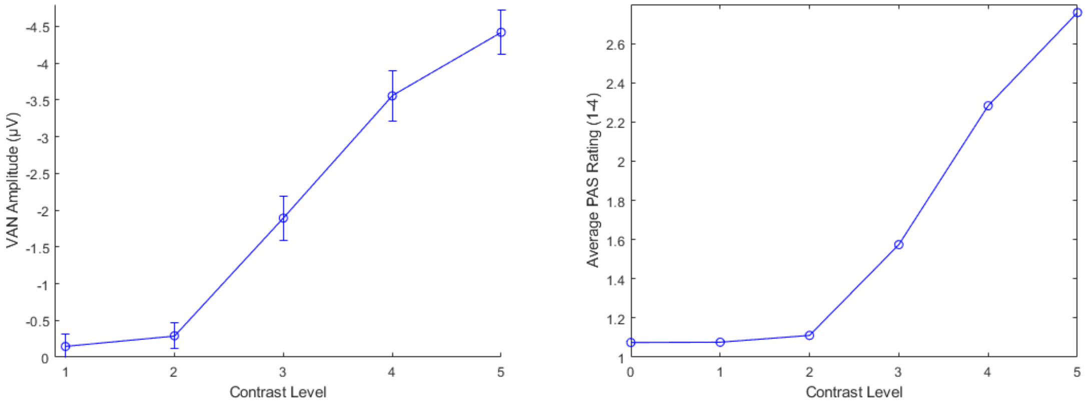

Figure 20. VAN amplitudes from the report pilots vs. average PAS ratings. Left: mean amplitudes of the VAN electrode pool recorded from the 10 pilot subjects. Right: average values given during the 4-point perceptual assessment scale rating system from the report condition in the final experiment.

If there was a significant increase from level 4 to 5, the overall pattern would be not sigmoidal but a linear trajectory. This mirrors the average value, 1-4, given by participants after each trial in the report condition to describe the visibility to each Gabor. Consider what the PAS scale ratings reflect—the overall strength and clarity of the stimulus as opposed to the binary of whether one consciously perceived it at all. The resemblance between these two curves prompts consideration of whether the VAN could possibly be indexing the *clarity* of the stimulus, a sort of perceptual resolution measurement, rather than conscious awareness itself. These two mechanisms are not completely detached; an increase in stimulus strength doesn't change the clarity if it is still too weak to be perceived. In this case, the VAN could still be considered an NCC, as it is always absent when stimuli are unseen and always present when stimuli are seen, but it’s amplitude profile above perceptual threshold might characterize subjective clarity which appears to show a different behavioral pattern than objective measures like orientation discrimination accuracy or binary measures of seen-vs-unseen.

This provokes the question of whether “seeing” a thing is truly binary, in the sense that all "unaware" stimuli fail to evoke neuronal activity and all "aware" stimuli evoke the same activity with the same strength. It follows that a signal related to perceptual awareness would not be evoked by a stimulus too weak to be perceived (in this case levels 1 and 2), but is it presumptuous to believe there would also cease to be significant differences in the signal strength after the threshold of visibility has been crossed? Perhaps, rather than a bifurcated dynamic, it is more reasonable to expect that NCCs will fluctuate in signal strength in variable ways in tandem with the continual changes in the nature of the experience as a sensory stimulus increases in strength. The binary of conscious awareness could be simply the first step of a broader picture; we may benefit from shifting our construction away from "not perceived vs. perceived" and toward something more like "not perceived vs. the entire spectrum of different qualities of perceptual experience." In this case, it doesn't matter if the VAN (or other signals, for that matter) continues to increase past threshold, as long as it is inactive without a consciously perceived stimulus and becomes active when the visibility threshold is crossed, and most importantly, the way that it changes above threshold should match some behavioral measure of the perceptual experience (such as the patterns of ratings using the PAS).

4.2. Rerouting the Investigation of the Fronto-Central N2

4.2.1. The fcN2 is not elicited by the experimentally tested parameters

This experimental paradigm did not replicate the results of Cohen et al. (2024) in producing a bifurcated fcN2 in the 250-300ms range. As this was the first study aimed specifically at isolating this ERP, the failure of the signal to appear was unexpected but may still illuminate insight into its nature. If the fronto-central N2 does not appear with these parameters–stimulus degradation with dynamic noise, a simple red dot probe distractor task, and remarkably simple Gabor stimuli, then what does prompt its appearance? The answer to this question may enable a better understanding of what it is actually indexing.

4.2.2. A retroactive investigative approach

This line of inquiry prompted a literature review of similar EEG primary research studies to determine which previous experiments may have unintentionally elicited an fcN2 signal. Twenty studies were compiled which, akin to this experiment, used the following paradigm measures:

1. Measurement of ERPs of conscious awareness using presentation of a stimulus
2. Manipulation of awareness of the stimulus using a controlled measure
3. A no-report segment coupled with a task to prevent or mitigate elicitation of the P3 signal during presentation of the stimulus.

Investigating available EEG data from these studies, 8 of them clearly indicated an fcN2-esque event, while 9 of them did not. Parsing these studies in an effort to illuminate the variables which prevent or induce the fcN2 can be broken down into three potential variables based on (1) the type of stimulus used, (2) the type of awareness manipulation, and (3) the type of task in the no-report condition. Observing these variables from a combinatorial and qualitative perspective reveals a few paths to insight.

4.2.3. Variables of Interest

The fcN2 is elicited by stimuli which are content-rich and specific

Examining the type of stimuli used to elicit awareness ERPs in the studies, a spectrum emerged which correlates the appearance of the fcN2 with certain types of stimuli. On one end of the spectrum are basic, geometric shapes and colors–squares, rectangles, Gabor gratings, blobs of color, etc. Of the five candidate studies which used these basic, vague stimuli and for which the presence of the fcN2 was measurable, none of them elicited the fcN2. This includes our study.

On the other end of the spectrum are specific objects such as animals, houses, objects and pictograms. In the six candidate studies which used these, five elicited the fcN2. Between these two ends are semantic stimuli, such as words, letters and vowel sounds, used by three of the candidate studies. Among them, one elicited the fcN2 and one did not, while the third is unable to be evaluated using available data (Hutchinson et al., 2024). The one notable outlying element in this scale is faces–only 4/8 studies which used facial stimuli yielded an fcN2; the number reduces to ¼ when observing studies that only used faces without other object-stimuli.

These findings suggest most confidently that basic, vague stimuli such as squares or gratings will routinely fail to yield an fcN2 signal, implying it is perhaps necessary for the stimulus to have some level of specificity or "richness" of content. With that said, while the richness may be a necessary factor, it decidedly does not guarantee elicitation. Furthermore, while intuition would suggest that human faces are more perceptually exciting as stimuli than houses or objects, there are markedly fewer appearances of the fcN2 from the former. Altogether, it seems that while rich stimuli are a helpful (and possibly necessary) marker of the fcN2's elusive activation, there are other variables which must be considered before the final outcome can be predicted.

Awareness manipulation methods vary greatly in efficacy of fcN2 elicitation

There were a variety of methods used across the studies to control the extent to which participants could or could not see the stimuli in varying conditions. Half of the twenty studies used some variation of an inattentional blindness paradigm. Among these, only two of them clearly elicited an fcN2, though 2 were unclear. This 25% appearance rate suggests that an IB paradigm is a generally ineffective means of eliciting the fcN2, especially compared to a method such as masking, which elicited the signal 3 out of 4 times. Notably, the only masking outlier was Pitts et al. (2014), which used very basic colors and shapes as stimuli. Given that masking was the central element of the Cohen et al. (2024) study which was changed for this follow-up, there is a compelling case to be made that it could boost the probability of the fcN2 appearing, even if it is not independently capable of producing it.

The No-Report condition: perceptual competition and cognitive involvement

Perhaps most compelling is the nature of the no-report condition wherein the EEG data itself is collected. Each study utilized a distractor task as a means of preventing attentional and task-related signals (namely the P3) from drowning out true NCCs, but the nature of this task was variable. Thirteen of the tasks used simple probe detections–that is, all the participants had to do was wait for an observable sensory input and then press a button. No further cognitive effort was required, nor did the distractor input change occur concurrently with the stimuli. Among the nine for which it was measurable, none of these studies elicited the fcN2.

The rest of the tasks possessed one or several of a qualifying set of variables I will label the "3 Cs": Concurrent, Competitive, and/or Cognitively involved. Examples include not only detecting a probe but determining if it fits a particular pattern (Colored Circles IB triangulation), counting the number of probe appearances (Masking long/short SOA), or engaging with a complex sensory experience such as a video game (Video game IB). It's arguable that a sensorily cross-modal distraction, such as an auditory task with a visual target stimulus, falls under cognitive involvement as well–the one study which utilized this did elicit the fcN2 (Auditory ID). Given that no study has utilized a 3C distractor task and failed to present an fcN2, the available data suggests that this method is an (as of yet) guaranteed means of generating the signal.

4.2.4. Combining the Variables and Drawing Conclusions

Permutations

While each of these three variables present compelling trends, it is necessary to consider them in combination with one another–after all, each study involved all three elements in some form. The table below provides a breakdown of whether the fcN2 is elicited in all eight possible combinations of the candidate conditions–rich stimuli, masking, and a 3C distractor task.

Rich Stimulus	**No**	**Yes**	**No**	**No**	**Yes**	**Yes**	**No**	**Yes**
Masking	No	No	Yes	No	Yes	No	Yes	Yes
3C No-Report Task	No	No	No	Yes	No	Yes	Yes	Yes
	⇩	⇩	⇩	⇩	⇩	⇩	⇩	⇩
FcN2?	No	?	No	?	Yes	Yes	?	Yes
Studies	1, 2	4, 3?*	5	6?**	7*** 8	9, 10, 11?	None	12, 11?, 3?

Table 1. Results of Experimental Variations on the appearance of the fcN2. Overview of nine combinations of the three variables of interest in fcN2 evocation. Studies are referenced using numeric keys listed below. Studies marked by a "?" are only arguably applicable to this set of conditions as one or more elements of their paradigm could be characterize as either containing or not containing a given element. For example, Sergent et al. (2021) used vowel sounds, which may or may not be considered "rich" stimuli. If the only studies reflecting a given permutation are labeled "?," so is the N2 outcome, even if the studies clearly did or did not measure it. Points marked by an asterisk warrant further elaboration. *Dellert et al., 2022 measured an fcN2, while DCF (triangulation) did *not*. ** Do auditory vowel stimuli count as specific? It is unclear if the results would change if the stimuli were shifted to something like houses, or at the other end, basic shapes. ***Masking SOA (triangulation) included faces but not objects, animals or houses.

Below is a list of the studies numbered in this table, labeled with a short-hand title and reference.

1. Twitchy Lines (shapes) (Pitts et al., 2012)
2. Twitchy Lines (left/right) (A. M. Harris et al., 2020)
3. Task-Uncertain Attentional Blindness (Dellert et al., 2021)
4. Dichoptic Color Fusion (Sarig et al., 2024)
5. Twitchy Masking (Pitts, Metzler, et al., 2014)
6. Auditory Bifurcation (Sergent et al., 2021)
7. Masking Stimulus Onset Asynchrony (Sarig et al., 2024)
8. Masking Bifurcation (Cohen et al., 2023)
9. Colored Circles IB (Sarig et al., 2024)
10. Video Game IB (Melloni et al., 2023)
11. Auditory ID (Schlossmacher et al., 2021)
12. Masking long/short Stimulus Onset Asynchrony (Cohen et al., 2020)

Corroborating the importance of these conditions, no study which used none of them has seen an fcN2 present. It is yet unclear if rich stimuli alone can elicit the N2, but clear that masking alone cannot. A 3C distractor *may* elicit the signal in situations where it otherwise wouldn't appear–it has yet to be tested in isolation with stimuli that are very clearly "vague," but simple vowel sounds arguably fit this category. A 3C distractor combined with rich stimuli does elicit the signal even in the absence of masking. Interestingly, masking combined with rich stimuli also generates the signal even in the absence of an involved distractor task–suggesting it is not strictly necessary for fcN2 appearance. The only permutation which lacks a study that could even arguably fit its conditions is a 3C distractor combined with masking in the absence of a content-rich stimulus such as a geometric shape. As the fcN2 has never been yielded from a study with vague stimuli, this could be an important final test condition to establish what is strictly necessary for elicitation.

Drawing Conclusions

The use of a no-report distractor task which utilizes all or many of the 3Cs seems to be the greatest tool presently available for isolating the fcN2, as it has not yet failed in doing so. Nonetheless, it is not strictly necessary, and the use of content-specific stimuli combined with a sensorily involved awareness manipulation such as masking can still generate the N2 without an involved distractor. It is unclear, however, if a 3C distractor can make up for a lack of specific stimuli, as the fcN2 has not yet been elicited in any study which does not use a content-rich stimulus. Nonetheless, the stimulus factor lacks the certainty power of the 3C distractor, as many studies which do use rich stimuli have still failed to evoke the N2. This offers two clear future directions for completing the permutations of these variables: using a 3C distractor with and without a masking method, both with a “bland” stimulus.

What does this actually mean for the fronto-central N2? While its initial discussion following Cohen et al. (2024) offered promise as a novel correlate with great theoretical significance, subsequent research has revealed a degree of elusivity which muddies this interpretation. If it is a broad NCC related to general awareness, it must be extraordinarily easy to drown out, in which case the best hope for parsing it is to continually pursue methodological adaptations that increase the clarity of ERP signal isolation. More likely, though, this ERP indexes a more complex processing event. Given that its markers–masking, rich stimuli, and cognitively complex/competitive/concurrent distractors–are all more sensorily involved and active than their alternatives, it’s possible that the fronto-central N2 demarcates something related to perception and filtration of stimuli in a cognitively competitive microtemporal environment. That is to say, when multiple stimuli are competing for attention, or when various cognitive mechanisms are concurrently active, the fcN2 occurs as the specific mechanistic route the brain must utilize at that time to achieve perceptual awareness of a stimulus.

It is also possible that the fcN2 is not even a specific ERP *per-se*, but rather the combined result of various neighboring charged activity events overlapping. In that case, it would explain why isolating the signal *too* well (by eliminating all distractions, extraneous ERPs and content-specific signals) could actually prevent it from appearing entirely–perhaps those distractions and extraneous signals are the very events which compose the fcN2 as an emergent pattern in the first place.

4.2.5. The fcN2 in the broader ERP ecosystem

A final point of characterization–less uniquely informative but readily apparent by observing the literature data–is that the fcN2 never appears in the absence of the newly discovered LVAP. It does *seem* to appear without the VAN, though labeling studies as "no-VAN" is somewhat misleading as it refers to the masking experiments in which the signal couldn't be isolated due to its residual latency shifts.

4.3. Novel Signals–the LVAP and LFN

4.3.1. Elusive Signal Patterns

The nature of the Late Visual Awareness Positivity (LVAP) and Late Frontal Negativity (LFN), particularly the former, prompt further investigation–both with the available data and in future experiments. The LVAP in particular, over the course of the experiment, developed a much broader scalp topography pattern than what was observed among the pilots. It appeared that there were multiple sources of positivity appearing in the posterior region of the scalp—one on each side laterally. Initial measurements of the grand average data using the pre-determined pool of electrodes in the late time window yielded an odd, exponential-looking curve pattern. This would imply that the underlying neural mechanisms become not only more active with each increase in contrast, but also that the rate of amplitude change increases with each step.

Dissecting the LVAP

Between this odd curve pattern and the observation of a second source of positivity in the grand average data, it felt prudent to plot each subject's scalp distributions individually during the time window of this ERP event. Interestingly, a very distinct pattern of inter-subject variability was revealed, with positivities appearing on the left and/or the right seemingly independently. Subject-by-subject observation of the most positive areas during the time window were used to create new electrode pools for the left and ride sides individually, although even when analyzed separately, each side maintained roughly the same expinential curve pattern as the combined LVAP, as well as roughly the same dual peak pattern within the time window.

It remains unclear whether the LVAP itself is a single cortical event or if this experiment elicited the appearance of multiple concurrent posterior positivities. While a more granular inspection of each side of the signal is unfortunately outside the scope of this thesis, it could be prudent to investigate separate time windows for the two sides as well, as the right positivity appears to be sustained past 650ms when the data is examined without a time window constraint.

4.3.2. Making Sense of the Late Signal System

It could be of interest to consider whether the Late Frontal Negativity is actually its own ERP, or if these events demarcate two sides of a single dipolar ignition event. Both appear to have the same pattern of amplitude change, which is supported by statistical analysis, and they occur within the same (fairly unique) time window. This is further corroborate by the left and right versions of the LVAP both presenting this amplitude pattern. There is significantly less distinct inter-subject variability in the LFN than in the LVAP, but perhaps the substantial spatial variation is the result of slight orientation changes at the other end of the dipole? Further experimentation would be necessary to determine if the two ERPs always behave concurrently in terms of activation and strength.

Regardless of whether these signals index the same circuitry activation, it appears that they are in some capacity related to awareness and perception. As they do not bifurcate (at least with the current limited electrode pools used for analysis), it is unlikely that they are true correlates of the awareness binary as we conceive it, although like the VAN, they remain completely inactive and unchanging when the contrast is too low to induce conscious awareness. Given this, it is at least unlikely that they relate to pure subconscious sensory processing a la the P1, as in that case one would expect amplitude changes between contrast levels 1 and 2.

Ultimately, the most that can be intuited from these events given this data is that they may reflect a common awareness-related process which occurs later than most other events and is variable in terms of its spatial distribution across subjects (Figure 19). Perhaps the LVAP is an ignition event which manifests over a particularly broad network of high-order cortical neurons why varies n anatomical distribution and firing pattern across the cortex; these neuronal characteristics are hallmarks of the theorized Global Neuronal Workspace. A curious follow-up experiment would be to introduce subsequent contrast levels above 5, perhaps 6-10, to determine how the amplitudes of these signals continue to change with even stronger visual stimuli.

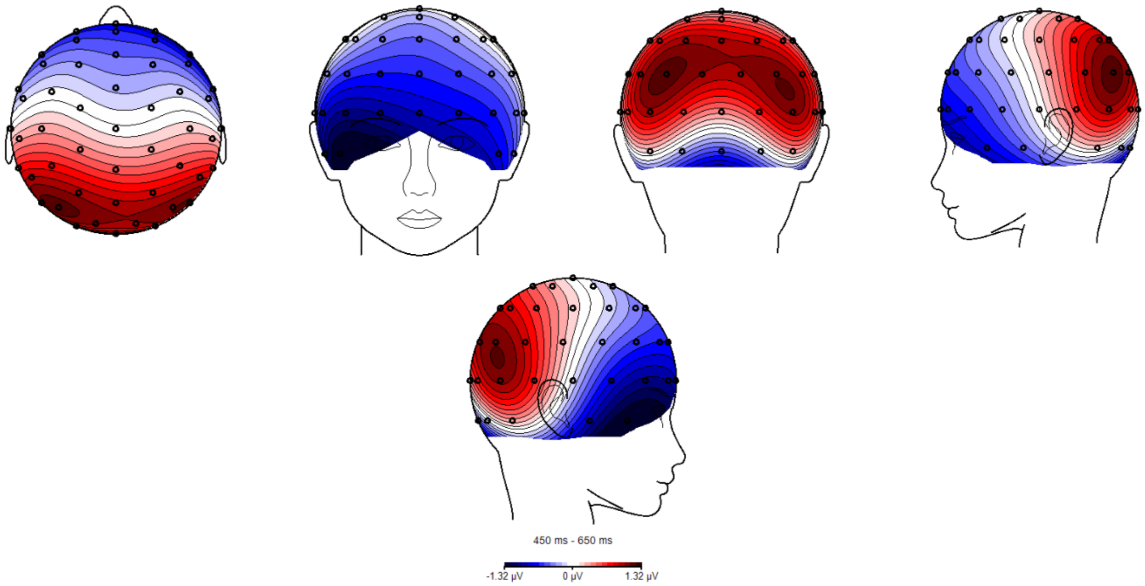

Figure 21. Five views of the 450-650ms window.
Grand average (n=30) scalp topography of activity in the late time window at the highest contrast level.

4.4. Conclusion

4.4.1. Contributions to the Search for the Neural Correlates of Consciousness

Throughout the collection and analysis of this data, insights have been gained into the nature of nearly half a dozen distinct ERPs. Even findings as simple as the successful attenuation of the P3b in the no-report condition supports and adds to the growing body of evidence that this common activity event corresponds to attention and the task of reporting itself. Establishing a thorough basis for this claim allows for firmer ground during theoretical elaboration on the relationship between this type of processing and "true" awareness. This study lends credence to the idea that there is a neurally founded distinction that can be made, although the extent of their independence will be the subject of future investigations.

The earliest occurring signal to be analyzed is also the one which has seen the most substantial empirical illumination. The Visual Awareness Negativity seems more closely associated with conscious awareness than ever, and this study weights against the idea that it corresponds with sensory input strength independent of perception, as it remains inactive until a clear catalyst point when conscious awareness begins as well. Furthermore, whether it bifurcates or corresponds more closely to something like the average PAS ratings, its pattern gives insight into both the empirical effects of the no-report paradigm and, potentially, the nature of awareness-associated neural amplitudes.

The lack of fronto-central N2 has offered an avenue for thorough investigation into what conditions prompt this signal to exist. It appears to be a surpisingly complex phenomenon corresponding to some combination of the richness of visual stimulus, and the concurrence of multiple sensory events, with cognitive competitive activity possibly playing a role as well. In any case, while it does not appear universally associated with conscious awareness, its time window and the fascinating conditions necessary to evoke it prompt further investigation into what this activity actually *means*.

Finally, the late signal window is rife with potential for further experimentation and isolation. These signals each present relatively similar patterns which neither bifurcate nor increase linearly with contrast strength; their curves imply a relationship to visual awareness that is not so simple as to remain static once awareness is achieved. In tandem with the VAN's behavior at higher contrast levels, these signals prompt theoretical discussion of what one can expect these amplitudes to become as input increases long after the awareness threshold. The spatial variety of the LVAP could be an indication of varying anatomical distribution of workspace neurons between individuals; perhaps it is even more fluid than subject-by-subject difference.

4.4.2. Limitations and Future Directions

Two key limitations of this design paradigm are made clear by the lack of clarity as to what happens with signals such as the VAN after threshold. The lack of insight into how these signal amplitudes change beyond threshold, as well as the potential movement of threshold itself both warrant future experiments in which there is a 6th, 7th, 8th, etc. contrast level that can be plotted and added to the amplitude curve. There are also the insights gleaned from investigating the conditions that evoke the fcN2. If this signal is to be pursued further, it seems that the extremely simple Gabor stimuli used in this study should be forgone for something more content-specific, such as an animal (or the faces used in the design that inspired this pursuit). While the implementation of a no-report task more cognitively engaging or concurrent than red dots which never overlap with the stimuli would be an as-of-yet guaranteed means of fcN2 evocation, it is also worth attempting to determine if rich stimuli alone could be sufficient. Once its precise conditions are understood, the question of what neurological process it actually indexes will likely become much more answerable. Finally, a follow-up experiment which re-implements the report condition could offer great insight into how the observable signals change in behavior when their strength is amplified greatly. This would be another means of remedying the hypothesized threshold shift in the VAN and detemining if this phenomenon has been less obviously affecting other ERP readings in the no-report condition as well.

Ultimately, these three key signals each offer insight into the circuitry which is responsible for generating the experience of conscious awareness. Overcoming the limitations of this study with directed follow-up paradigms will undoubtedly provide future insights into what circuitry and activity in the human brain underlies the very existence of subjective experience.

Appendix A: Full Dataset

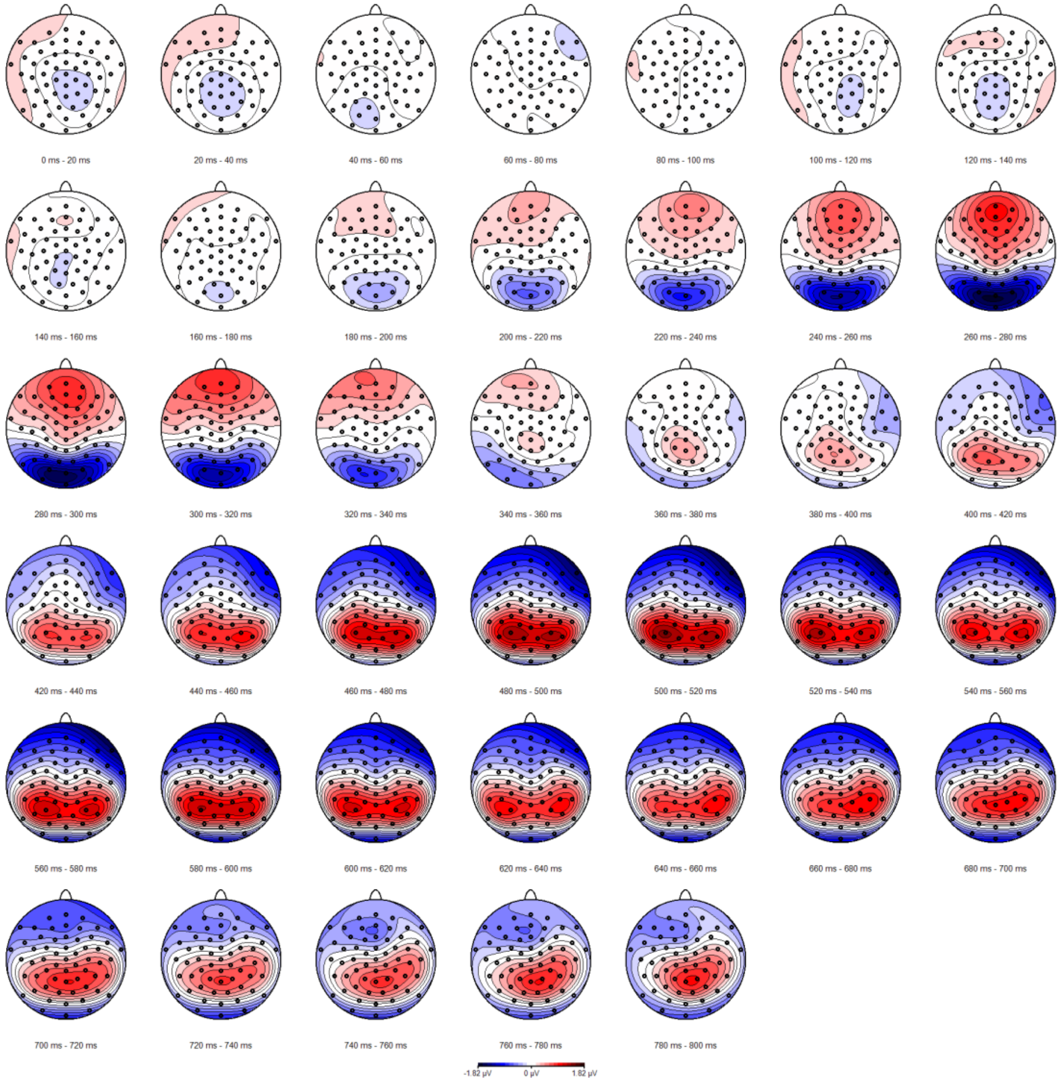

Supplement 1. Scalp topographies from 0-800ms.

Grand average of 30 subject EEG recordings at contrast level 5 over the full standard ERP interval, beginning with stimulus onset at 0ms and ending 0.8 seconds afterward.

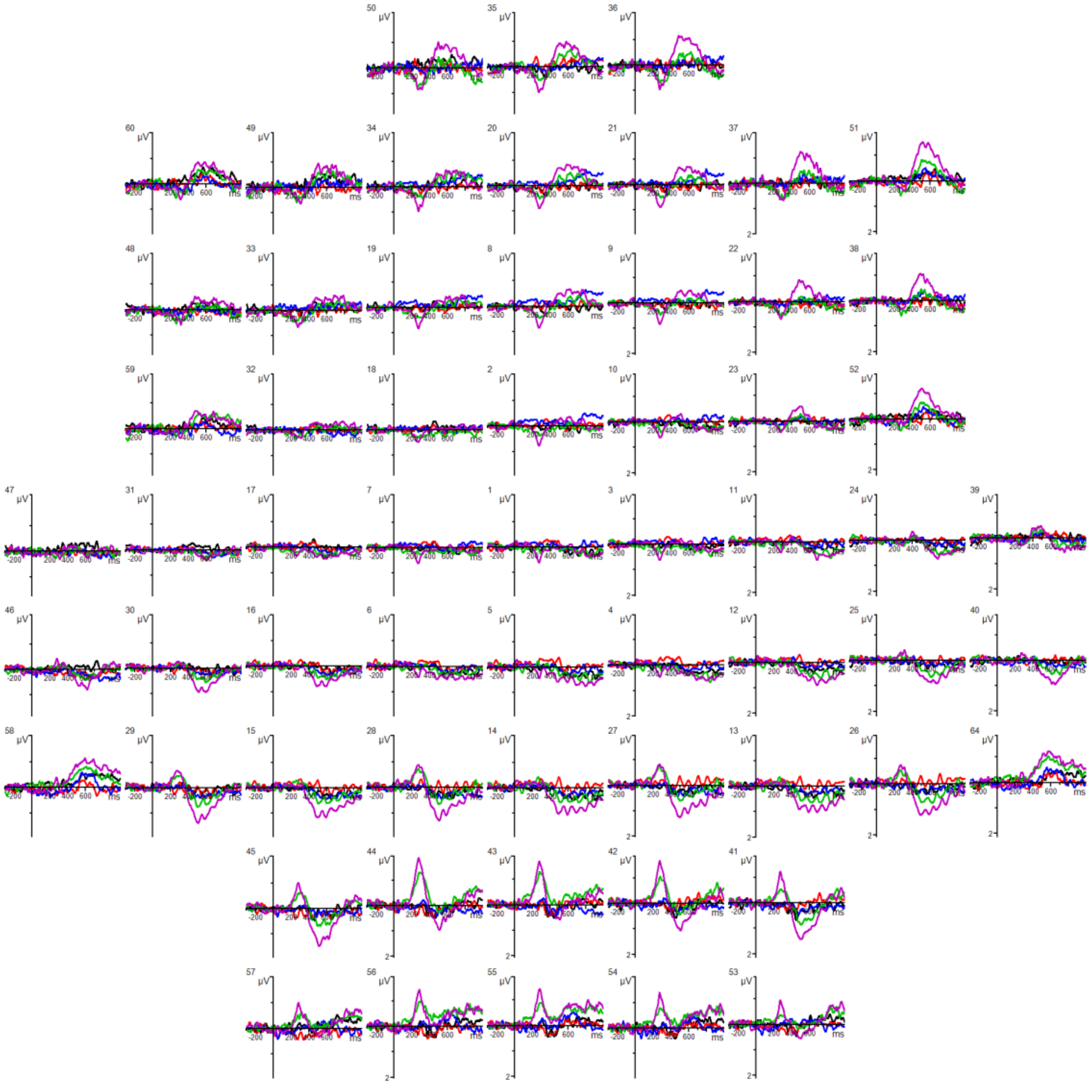

Supplement 2. Amplitudes of each electrode from 0-800ms.

Full waveform data of each electrode, grand averaged (n = 30), across the full 0-800ms time window for each contrast level (black = 1, red = 2, blue = 3, green = 4, purple = 5).

Bibliography

Alexander, G. E., DeLong, M. R., & Strick, P. L. (1986). Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annual Review of Neuroscience*, *9*, 357–381. https://doi.org/10.1146/annurev.ne.09.030186.002041

Bachmann, T., & Hudetz, A. G. (2014). It is time to combine the two main traditions in the research on the neural correlates of consciousness: C = L × D. *Frontiers in Psychology*, *5*, 940. https://doi.org/10.3389/fpsyg.2014.00940

Berkovitch, L., Dehaene, S., & Gaillard, R. (2017). Disruption of Conscious Access in Schizophrenia. *Trends in Cognitive Sciences*, *21*(11), 878–892. https://doi.org/10.1016/j.tics.2017.08.006

Block, N. (2005). Two neural correlates of consciousness. *Trends in Cognitive Sciences*, *9*(2), 46–52. https://doi.org/10.1016/j.tics.2004.12.006

Britton, J. W., Frey, L. C., Hopp, J. L., Korb, P., Koubeissi, M. Z., Lievens, W. E., Pestana-Knight, E. M., & St. Louis, E. K. (2016). *Electroencephalography (EEG): An Introductory Text and Atlas of Normal and Abnormal Findings in*

Adults, Children, and Infants (E. K. St. Louis & L. C. Frey, Eds.). American Epilepsy Society. http://www.ncbi.nlm.nih.gov/books/NBK390354/

Butti, C., Santos, M., Uppal, N., & Hof, P. R. (2013). Von Economo neurons: Clinical and evolutionary perspectives. *Cortex; a Journal Devoted to the Study of the Nervous System and Behavior*, *49*(1), 312–326. https://doi.org/10.1016/j.cortex.2011.10.004

Carrasco, M. (2018). How visual spatial attention alters perception. *Cognitive Processing*, *19*(Suppl 1), 77–88. https://doi.org/10.1007/s10339-018-0883-4

Cohen, M. A., Dembski, C., Ortego, K., Steinhibler, C., & Pitts, M. (2023). *Neural signatures of visual awareness independent of post-perceptual processing* (p. 2023.06.09.543951). bioRxiv. https://doi.org/10.1101/2023.06.09.543951

Cohen, M. A., Ortego, K., Kyroudis, A., & Pitts, M. (2020). Distinguishing the Neural Correlates of Perceptual Awareness and Postperceptual Processing. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *40*(25), 4925–4935. https://doi.org/10.1523/JNEUROSCI.0120-20.2020

Dehaene, S., Lau, H., & Kouider, S. (2017). What is consciousness, and could machines have it? *Science (New York, N.Y.)*, *358*(6362), 486–492. https://doi.org/10.1126/science.aan8871

Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition*, *79*(1–2), 1–37. https://doi.org/10.1016/s0010-0277(00)00123-2

Del Cul, A., Baillet, S., & Dehaene, S. (2007). Brain dynamics underlying the nonlinear threshold for access to consciousness. *PLoS Biology*, *5*(10), e260. https://doi.org/10.1371/journal.pbio.0050260

Dellert, T., Müller-Bardorff, M., Schlossmacher, I., Pitts, M., Hofmann, D., Bruchmann, M., & Straube, T. (2021). Dissociating the Neural Correlates of Consciousness and Task Relevance in Face Perception Using Simultaneous EEG-fMRI. *The Journal of Neuroscience*, *41*(37), 7864–7875. https://doi.org/10.1523/JNEUROSCI.2799-20.2021

Dembski, C., Koch, C., & Pitts, M. (2021). Perceptual awareness negativity: A physiological correlate of sensory consciousness. *Trends in Cognitive Sciences*, *25*(8), 660–670. https://doi.org/10.1016/j.tics.2021.05.009

Eklund, R., & Wiens, S. (2018). Visual awareness negativity is an early neural correlate of awareness: A preregistered study with two Gabor sizes. *Cognitive, Affective & Behavioral Neuroscience*, *18*(1), 176–188. https://doi.org/10.3758/s13415-018-0562-z

Fang, Z., Dang, Y., Ping, A., Wang, C., Zhao, Q., Zhao, H., Li, X., & Zhang, M. (2025). Human high-order thalamic nuclei gate conscious perception through the thalamofrontal loop. *Science*, *388*(6742), eadr3675. https://doi.org/10.1126/science.adr3675

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

Frässle, S., Sommer, J., Jansen, A., Naber, M., & Einhäuser, W. (2014). Binocular rivalry: Frontal activity relates to introspection and action but not to perception. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *34*(5), 1738–1747. https://doi.org/10.1523/JNEUROSCI.4403-13.2014

Harris, A. M., Dux, P. E., & Mattingley, J. B. (2020). Awareness is related to reduced post-stimulus alpha power: A no-report inattentional blindness study. *European Journal of Neuroscience*, *52*(11), 4411–4422. https://doi.org/10.1111/ejn.13947

Harris, K. D., & Shepherd, G. M. G. (2015). The neocortical circuit: Themes and variations. *Nature Neuroscience*, *18*(2), 170–181. https://doi.org/10.1038/nn.3917

Hatamimajoumerd, E., Ratan Murty, N. A., Pitts, M., & Cohen, M. A. (2022). Decoding perceptual awareness across the brain with a no-report fMRI masking paradigm. *Current Biology: CB*, *32*(19), 4139-4149.e4. https://doi.org/10.1016/j.cub.2022.07.068

Kapoor, V., Dwarakanath, A., Safavi, S., Werner, J., Besserve, M., Panagiotaropoulos, T. I., & Logothetis, N. K. (2022). Decoding internally generated transitions of conscious contents in the prefrontal cortex without subjective reports. *Nature Communications*, *13*(1), 1535. https://doi.org/10.1038/s41467-022-28897-2

Koch, C., Massimini, M., Boly, M., & Tononi, G. (2016). Neural correlates of consciousness: Progress and problems. *Nature Reviews. Neuroscience*, *17*(5), 307–321. https://doi.org/10.1038/nrn.2016.22

Kouider, S., Stahlhut, C., Gelskov, S. V., Barbosa, L. S., Dutat, M., de Gardelle, V., Christophe, A., Dehaene, S., & Dehaene-Lambertz, G. (2013). A neural marker of perceptual consciousness in infants. *Science (New York, N.Y.)*, *340*(6130), 376–380. https://doi.org/10.1126/science.1232509

Lamme, V. A. F. (2006). Towards a true neural stance on consciousness. *Trends in Cognitive Sciences*, *10*(11), 494–501. https://doi.org/10.1016/j.tics.2006.09.001

Lamme, V. A. F. (2018). Challenges for theories of consciousness: Seeing or knowing, the missing ingredient and how to deal with panpsychism. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *373*(1755), 20170344. https://doi.org/10.1098/rstb.2017.0344

Larkum, M. (2013). A cellular mechanism for cortical associations: An organizing principle for the cerebral cortex. *Trends in Neurosciences*, *36*(3), 141–151. https://doi.org/10.1016/j.tins.2012.11.006

Marti, S., & Dehaene, S. (2017). Discrete and continuous mechanisms of temporal selection in rapid visual streams. *Nature Communications*, *8*(1), 1955. https://doi.org/10.1038/s41467-017-02079-x

Marti, S., King, J.-R., & Dehaene, S. (2015). Time-Resolved Decoding of Two Processing Chains during Dual-Task Interference. *Neuron*, *88*(6), 1297–1307. https://doi.org/10.1016/j.neuron.2015.10.040

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

Melloni, L., Mudrik, L., Pitts, M., Bendtz, K., Ferrante, O., Gorska, U., Hirschhorn, R., Khalaf, A., Kozma, C., Lepauvre, A., Liu, L., Mazumder, D., Richter, D., Zhou, H., Blumenfeld, H., Boly, M., Chalmers, D. J., Devore, S., Fallon, F., … Tononi, G. (2023). An adversarial collaboration protocol for testing contrasting predictions of global neuronal workspace and integrated information theory. *PLOS ONE*, *18*(2), e0268577. https://doi.org/10.1371/journal.pone.0268577

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

Biological Sciences, *373*(1755), 20170357. https://doi.org/10.1098/rstb.2017.0357

Odegaard, B., Knight, R. T., & Lau, H. (2017). Should a Few Null Findings Falsify Prefrontal Theories of Conscious Perception? *Journal of Neuroscience*, *37*(40), 9593–9602. https://doi.org/10.1523/JNEUROSCI.3217-16.2017

Parvizi, J., & Damasio, A. (2001). Consciousness and the brainstem. *Cognition*, *79*(1–2), 135–160. https://doi.org/10.1016/s0010-0277(00)00127-x

Pitts, M. A., Lutsyshyna, L. A., & Hillyard, S. A. (2018). The relationship between attention and consciousness: An expanded taxonomy and implications for "no-report" paradigms. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *373*(1755), 20170348. https://doi.org/10.1098/rstb.2017.0348

Pitts, M. A., Martínez, A., & Hillyard, S. A. (2012). Visual processing of contour patterns under conditions of inattentional blindness. *Journal of Cognitive Neuroscience*, *24*(2), 287–303. https://doi.org/10.1162/jocn_a_00111

Pitts, M. A., Metzler, S., & Hillyard, S. A. (2014). Isolating neural correlates of conscious perception from neural correlates of reporting one's perception. *Frontiers in Psychology*, *5*, 1078. https://doi.org/10.3389/fpsyg.2014.01078

Pitts, M. A., Padwal, J., Fennelly, D., Martínez, A., & Hillyard, S. A. (2014). Gamma band activity and the P3 reflect post-perceptual processes, not visual awareness. *NeuroImage*, *101*, 337–350. https://doi.org/10.1016/j.neuroimage.2014.07.024

Sakata, S., & Harris, K. D. (2009). Laminar structure of spontaneous and sensory-evoked population activity in auditory cortex. *Neuron*, *64*(3), 404–418. https://doi.org/10.1016/j.neuron.2009.09.020

Sarig, A., Liad Mudrik, Pitts, M., & Schurger, A. (2024). *Triangulating Consciousness*. OSF Registries. https://doi.org/10.17605/OSF.IO/TFPE4

Schelonka, K., Graulty, C., Canseco-Gonzalez, E., & Pitts, M. A. (2017). ERP signatures of conscious and unconscious word and letter perception in an inattentional blindness paradigm. *Consciousness and Cognition*, *54*, 56–71. https://doi.org/10.1016/j.concog.2017.04.009

Schlossmacher, I., Dellert, T., Bruchmann, M., & Straube, T. (2021). Dissociating neural correlates of consciousness and task relevance during auditory processing. *NeuroImage*, *228*, 117712. https://doi.org/10.1016/j.neuroimage.2020.117712

Self, M. W., Kooijmans, R. N., Supèr, H., Lamme, V. A., & Roelfsema, P. R. (2012). Different glutamate receptors convey feedforward and recurrent processing in macaque V1. *Proceedings of the National Academy of Sciences of the United States of America*, *109*(27), 11031–11036. https://doi.org/10.1073/pnas.1119527109

Sergent, C., Baillet, S., & Dehaene, S. (2005). Timing of the brain events underlying access to consciousness during the attentional blink. *Nature Neuroscience*, *8*(10), 1391–1400. https://doi.org/10.1038/nn1549

Sergent, C., Corazzol, M., Labouret, G., Stockart, F., Wexler, M., King, J.-R., Meyniel, F., & Pressnitzer, D. (2021). Bifurcation in brain dynamics reveals a signature of conscious processing independent of report. *Nature Communications*, *12*(1), 1149. https://doi.org/10.1038/s41467-021-21393-z

Shafto, J. P., & Pitts, M. A. (2015). Neural Signatures of Conscious Face Perception in an Inattentional Blindness Paradigm. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *35*(31), 10940–10948. https://doi.org/10.1523/JNEUROSCI.0145-15.2015

Storm, J. F., Klink, P. C., Aru, J., Senn, W., Goebel, R., Pigorini, A., Avanzini, P., Vanduffel, W., Roelfsema, P. R., Massimini, M., Larkum, M. E., & Pennartz, C. M. A. (2024). An integrative, multiscale view on neural theories of consciousness. *Neuron*, *112*(10), 1531–1552. https://doi.org/10.1016/j.neuron.2024.02.004

Yu, J., Anderson, C. T., Kiritani, T., Sheets, P. L., Wokosin, D., Wood, L., & Shepherd, G. M. G. (2008). Local-circuit phenotypes of layer 5 neurons in motor-frontal cortex of YFP-H mice. *Frontiers in Neural Circuits*, *2*. https://doi.org/10.3389/neuro.04.006.2008

Zhang, S., Xu, M., Kamigaki, T., Do, J. P. H., Chang, W.-C., Jenvay, S., Miyamichi, K., Luo, L., & Dan, Y. (2014). Long-Range and Local Circuits for Top-Down Modulation of Visual Cortical Processing. *Science (New York, N.Y.)*, *345*(6197), 660–665. https://doi.org/10.1126/science.1254126