Decoding EEG Bifurcation Dynamics Around Visual Detection Threshold in No-Report

---

A Thesis

Presented to

The Established Interdisciplinary Committee for Neuroscience

Reed College

In Partial Fulfillment

of the Requirements for the Degree

Bachelor of Arts

Angelica Nicolacoudis

May 2025

Approved for the Committee

(Neuroscience)

_____

Michael Pitts

Acknowledgments

As cliché as it sounds, words cannot describe how thankful I am to Michael for introducing me to cognitive neuroscience, for helping build my confidence as a researcher, and for always believing in my ability to succeed. My experience in the SCALP Lab has been life changing, and I could not be more grateful.

I'd also like to thank my mom for her never-ending support, for teaching me to persevere and trust myself. Thank you for listening to me lament and perseverate and complain over the last 4 years – your patience and love got me through every challenge. Thank you to my sister for watching out for me, and for never missing an opportunity to explain a new way to find a derivative. Thank you to my aunt for your unconditional love, and for teaching me to always be curious about the world. Thank you to my grandpa for your unwavering support and belief in my success.

Thank you to everyone at the bookstore for your kindness, support, music and movie recommendations, silly pranks, and community. I cannot imagine what my Reed experience would have been like if it weren't for you.

Finally, thank you to my friends. You've made my time at Reed truly special. Thank you to Sophie for agreeing to be roommates freshman year, for all the adventures and coffee shop gossip sessions and uncontrollable laughs. Thank you to Angie, Cedar, Rowan, and Haley for being incredible friends and housemates, and for all the movie nights, impromptu kitchen chats, and walks for ice cream.

Also, shout out to everyone at Papaccino's for keeping me sane and caffeinated.

Table of Contents

Chapter 1: Introduction **1**

1.1 The Search for Neural Correlates of Consciousness 1

1.2 Perceptual Reports and the Rise of No-Report Paradigms 2

1.3 Bifurcation Dynamics of Perception 5

1.3.1 Limitations of Backwards Masking in Investigating Bifurcation Dynamics 9

1.4 Candidate NCCs 11

1.4.1 Event Related Potentials 12

Visual Awareness Negativity 12

Fronto-Central N2 13

1.4.2 Metastability 13

Multivariate Pattern Analysis 14

Temporal Generalization of Decoding 17

1.5 Sensory Versus Cognitive Theories of Consciousness 19

1.6 The Current Study 21

Chapter 2: Methods **23**

2.1 Participants 23

2.2 Stimuli 23

2.3 Design 25

2.3.1 Thresholding Procedure 25

2.3.2 Pilot Experiment 1 26

2.3.3 Pilot Experiment 2 30

2.3.4 Main Experiment 34
2.4 Analysis 37
2.4.1 Behavioral 37
2.4.2 Data Segmentation and Pre-Processing 37
2.4.3 Multivariate Pattern Analysis 38
Decoding 38
Temporal Generalization of Decoding 39
Chapter 3: Results 41
3.1 Behavioral Results 41
3.2 Decoding Stimulus Presence at Each Contrast Level 47
3.3 Seen vs. Unseen Decoding 53
3.4 Decoding Adjacent Contrast Levels 55
3.5 Decoding Stimulus Orientation at Each Contrast Level 60
Chapter 4: Discussion 63
4.1 Temporal Generalization as a Potential NCC 64
4.2 Theoretical Implications 70
4.3 Limitations and Future Directions 72
Bibliography 75

List of Figures

Figure 1: Behavioral results from Cohen et al. (2024) showing a bifurcated pattern of perception. 7

Figure 2: Example of same-time decoding of face stimulus visibility in an inattentional blindness paradigm. 16

Figure 3: Example temporal generalization matrices derived from differing neural generators (King & Dehaene, 2014). 18

Figure 4: Example of a "left" Gabor grating (135°). 24

Figure 5: Still example of continuous dynamic noise. 25

Figure 6: Average orientation discrimination task performance from the first pilot experiment. 28

Figure 7: Average percentage of Gabor gratings rated as seen during the first pilot experiment. 29

Figure 8: Topographic scalp distributions from the no-report and report conditions of the first pilot experiment. 30

Figure 9: Average orientation discrimination task performance from the second pilot experiment. 32

Figure 10: Average percentage of Gabor gratings rated as seen during the second pilot experiment. 33

Figure 11: Topographic scalp distributions from the no-report and report conditions of the second pilot experiment. 34

Figure 12: Design of the final stimulus degradation paradigm. 36

Figure 13: Behavioral functions from objective and subjective participant reports. 42

Figure 14: Distribution of PAS ratings at each contrast level. 44

Figure 15: Average PAS Rating (1-4) at each contrast level. 45

Figure 16: Linear and Non-Linear Regression Fitting of Behavioral Data. 46

Figure 17: Linear and exponential fitting of average PAS ratings. 47

Figure 18: Same-time decoding of stimulus presence at each contrast level. 49

Figure 19: Temporal generalization of stimulus presence decoding at each contrast level. 50

Figure 20: Linear and non-linear regression of average temporal generalization classifier accuracy for stimulus presence decoding at each contrast level. ... 52

Figure 21. Decoding stimulus visibility contrast level 5 from level 1 (**A**) and level 4 from level 2 (**B**). 54

Figure 22. Same-time decoding of adjacent contrast levels. 57

Figure 23. Temporal generalization of adjacent contrast level decoding. 59

Abstract

Studies aimed at isolating neural signals linked with conscious perception overwhelmingly utilize a binary contrastive approach, manipulating awareness to compare brain responses between conscious and unconscious conditions. This method, while informative, often yields numerous candidate signals – which can be further confounded by signals related to reporting. Seeking more precise isolation, Cohen et al. (2024) used a no-report visual masking paradigm in which stimulus visibility was parametrically manipulated such that perception was either below, at, or above perceptual threshold. Neural responses in the no-report condition were compared with the psychometric function of visibility obtained in the report condition. Decoding of stimulus presence revealed temporal generalization only in the report condition, complicating previous findings of metastability in no-report.

Here, we present a variation of this experiment, with simpler stimuli and no masks, in which visibility was manipulated by varying stimulus contrast across five levels (2 below, 1 at, 2 above threshold). Multivariate pattern analysis was used to compare behavioral functions to scalp EEG data, decoding stimulus features both along a same-time training-and-testing diagonal and at all other time points using the temporal generalization method. Temporal generalization of stimulus presence decoding showed significant spreading from 200-400ms and 400-600ms at the highest contrast level, with non-linear increases in classification accuracy with increasing visibility, matching closely with behavioral functions. Seen versus unseen temporal generalization revealed similar results, confirming the link between metastability and conscious awareness. These results support a

unified, two-stage theory of conscious processing in which both early and late metastable signals are required for awareness.

For Marmie. Your love is with me all the time, everywhere.

Chapter 1: Introduction

1.1 The Search for Neural Correlates of Consciousness

One of the major endeavors in cognitive neuroscience has been to identify the neural basis for consciousness. Consciousness is typically experienced as a series of varying states of awareness, such as the transition from wakefulness to drowsiness to sleep. The physiological basis for consciousness has long been situated in the brain: damage to the visual cortex leads to perception without conscious vision, or “blindsight” (Humphrey & Weiskrantz, 1967); intracranial stimulation of the fusiform face area produces illusory face-like visual experiences, or “facephenes” (Schalk et al., 2017); and intracranial recordings of brain activity reveal increased neural signaling in occipitotemporal, inferior parietal, and prefrontal cortex for consciously perceived versus unconsciously processed face stimuli (Herman et al., 2019). While understanding the neural mechanisms underlying states of consciousness (i.e. awake, asleep, comatose) is insightful, such states are background enabling conditions necessary but not sufficient for conscious experiences. In other words, even when in an awake, "conscious state," most processing in the brain remains unconscious, with only a subset of these processes leading to consciousness. In this context, consciousness can be defined as subjective experience, or "what it is like" to think, act, and perceive (Koch, 2016). The search for a neural basis of consciousness has largely centered around identifying neural correlates of consciousness (NCCs), or the neuronal events jointly necessary and sufficient for a conscious experience

(Dehaene, 2014). By measuring responses to specific stimuli across sensory modalities (i.e. auditory, visual, somatosensory, etc.), content-specific correlates of perceptual awareness can be identified. Comparing content-specific NCCs may reveal the neural substrate supporting subjective experience, finally answering the scientific and philosophical conundrum that is consciousness (Koch et al., 2016).

1.2 Perceptual Reports and the Rise of No-Report Paradigms

A plethora of strategies have been employed to manipulate perception and measure corresponding brain activity – inattentional blindness, dichoptic color fusion, rapid serial visual presentation, and backwards masking, to name a few – yet complications continue to render the search for NCCs highly debated (Koch et al., 2016). Separating signatures of consciousness from surrounding processes, for instance, has called some candidate NCCs into question, particularly the P3b (Cohen et al., 2020; Hatamimajoumerd et al., 2022). Typically, researchers rely on subject reports to determine when a stimulus was perceived or not. In a visual backwards masking study conducted by Del Cul et al. (2007), for example, participants were instructed to indicate, or report, whether they perceived a target number different than 5, then rate their subjective experience of the target number on a continuous scale. Subsequent neural responses to perceived versus not-perceived trials were collected using electroencephalography (EEG), revealing a significant positive component in central electrodes from 370-500ms, identified as the P3b. Several other studies have identified the P3b (characterized by late positivity in central electrodes from

~300-600ms) as a potential correlate of consciousness, all relying on trial-by-trial reports to assess awareness of their respective stimuli (Sergent et al., 2005; Kouider et al., 2013; Koivisto et al., 2008; Fernandez-Duque et al., 2003). In more recent years, though, the P3b has been shown to reflect post-perceptual processes, rather than perception alone (Gaillard et al., 2009; Panagiotaropoulus, 2024; Sergent et al., 2021; Pitts et al., 2014).

To better isolate signals associated with conscious perception and reduce post-perceptual effects, researchers have developed "no-report" paradigms in which subjects do not report on their perceptual experience or make judgements about target stimuli, and perception is instead inferred indirectly via separate report conditions and/or physiological measures such as eye movements, pupillometry, etc. (Frässle et al., 2014; Kapoor et al., 2022; Pitts et al., 2014; Tsuchiya et al., 2015; Block, 2019). Mounting evidence has shown that candidate NCCs like the P3b disappear in no-report paradigms, further highlighting the need to minimize the effects of post-perceptual processing (Cohen et al., 2024; Cohen et al., 2020; Pitts et al., 2012; Pitts et al., 2014; Shafto & Pitts, 2015).

Still, some have questioned the ability of no-report paradigms to truly eliminate the neural precursors and/or consequences of consciousness, which confound the isolation of true NCCs. Koch (2016) includes in his definition of NCCs the requisite that a candidate NCC must reflect the *minimal set* of neuronal events jointly necessary and sufficient for a conscious experience. The implementation of no-report paradigms follows from this requisite condition, promising to distinguish perceptual awareness from processes like introspection, working memory, and meta-cognition (Tsuchiya et al., 2015). As Duman et al. (2022) argue, though, consciousness may involve some of these surrounding processes. One process closely tied to consciousness is attention (Graziano & Kastner, 2011; Mack & Rock, 1998). Proponents of Attention Schema Theory

(AST) (discussed more thoroughly in section 1.5), for instance, suggest that attentional processes lead the brain to form a sense of subjective experience, inextricably linking attention with consciousness (Graziano & Webb, 2015). In fact, many studies take advantage of the apparent link between attention and consciousness, manipulating attention to render stimuli perceptually invisible, such as in the case of inattentional blindness (Mack and Rock, 1998; Pitts et al., 2014, Schelonka et al., 2017; Tsuchiya et al., 2015). Whereas report paradigms have been criticized for overestimating the NCC, no-report paradigms may fall victim to underestimating the NCC (Duman et al., 2022).

An additional criticism of no-report paradigms is best summed up by Block (2019)'s description of the "bored monkey" problem. In visual no-report paradigms, subjects are often required to passively view target stimuli, maintaining attention while abstaining from explicit report. Under these conditions, Block (2019) argues, participants likely still think, wonder, question, or muse about the target stimulus, even when instructed to ignore or disregard them. The only way to fully eliminate post-perceptual cognitive processes from perceptual awareness, then, is to implement "no-cognition" or "no-differential-post-perceptual-processing" paradigms in which changes in perception are inconspicuous and do not draw the subject's attention (Block, 2019; Duman et al., 2022). Achieving such a paradigm has proven extremely difficult, although some studies have made progress (Brascamp et al., 2015).

Despite their apparent shortcomings, no-report paradigms still serve a vital role in identifying candidate NCCs. An increasing number of no-report paradigms have implemented auxiliary probe tasks to prevent, or at least reduce the effect of, the bored monkey problem (Cohen et al., 2020, 2024; Hesse & Tsao, 2020; Sergent et al., 2021; Dellert et al., 2022). These studies continue to show an absence of the P3b, as well as a limited role of prefrontal areas in awareness

(although this does not negate a possible role of prefrontal areas in consciousness, for further discussion see Panagiotaropoulus et al., 2024 and Duman et al., 2022). Overall, no-report paradigms reduce the chance of overestimating consciousness, allowing researchers to more precisely identify candidate NCCs, especially when paired with gradual manipulations of perception (see section 1.3 below).

1.3 Bifurcation Dynamics of Perception

NCC studies primarily rely on two types of perceptual manipulations: discrete or gradual. The former method involves the comparison of stimuli that are either never perceived or always perceived, such as with masking studies using a short (ex. 8ms) and long (ex. 80ms) stimulus onset asynchrony (SOA) (Cohen et al., 2020), or inattentional blindness studies that contrast attended stimuli with unattended stimuli (Pitts et al., 2014). While such studies are useful for identifying potential NCCs, information about the transition from non-unconscious to conscious processing may be lost and the neural mechanisms underlying consciousness overestimated. The latter method likely better captures the neural basis of perceptual transitions, relying on the manipulation of stimuli around perceptual thresholds (Sergent & Dehaene, 2004; Del Cul et al., 2007; Sergent et al., 2021; Cohen et al., 2024). Previous studies have shown that when stimulus strength increases in a linear manner, perception changes in an all-or-nothing fashion around perceptual threshold (Dehaene, 2014; Del Cul et al., 2007) following a standard sigmoid-shaped psychometric curve with a bimodal distribution separating conscious from unconscious processing (Helmholtz & Southall, 1924; Boring, 1950; Sergent et al., 2021) (Figure 1). In a behavioral attentional blink study conducted by Sergent and Dehaene (2004), participants

were instructed to rate their perception of a target stimulus (the word "five") using a continuous scale ("not seen" at the far left and "maximal visibility" at the far right, with 21 contiguous positions in between). Despite designing their subjective awareness scale to account for continuous changes in perception, participants overwhelmingly used the scale in an all-or-none fashion, with stimuli within the attentional blink (i.e. lag 3) rated equally at 0% or 100% visibility, and stimuli outside of the attentional blink (i.e. lag 8) rated almost entirely at 100% visibility. These findings demonstrate that perception displays qualitative bifurcation dynamics around perceptual threshold.

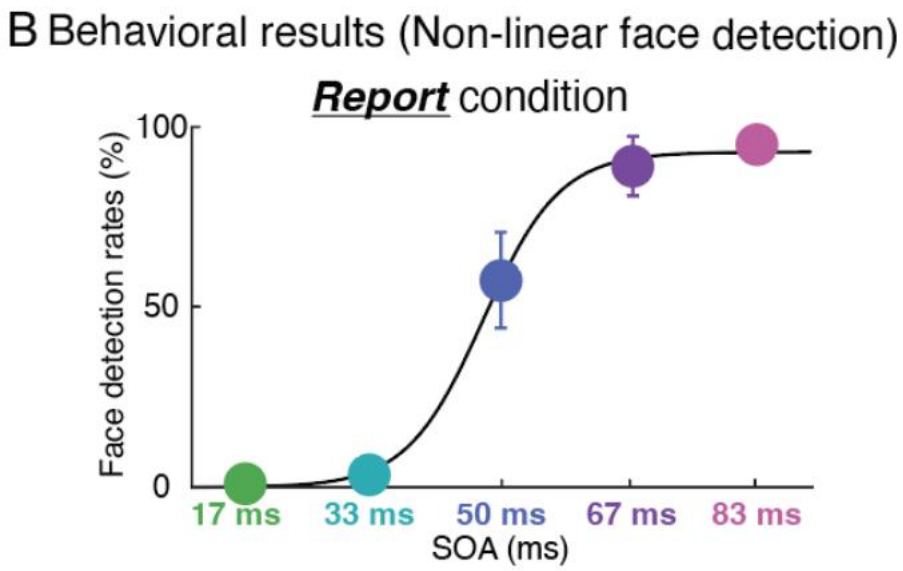

Figure 1: Behavioral results from Cohen et al. (2024) showing a bifurcated pattern of perception.

Cohen et al. (2024) first identified a common perceptual threshold across participants, then linearly increased the SOA to yield non-linear, bimodal behavioral responses. Behavioral data was acquired during the report condition, during which participants were instructed to respond whether they had seen a face or not. Accuracy on the face detection task is plotted here by SOA, demonstrating that at intervals, faces were almost never perceived, at threshold perceived around 50% of the time, and at longer intervals perceived almost every time. The resulting behavioral function resembles a classic sigmoid psychometric curve (Helmholtz & Southall, 1924; Boring, 1950). Cohen, M. A., Dembski, C., Ortego, K., Steinhibler, C., & Pitts, M. (2024). Neural signatures of visual awareness independent of postperceptual processing. Cerebral Cortex, 34(11), bhae415. https://doi.org/10.1093/cercor/bhae415.

A handful of subsequent studies have attempted to link behavioral bifurcation dynamics with neural responses. To better understand brain dynamics around perceptual threshold, Del Cul et al. (2007) recorded continuous EEG during a backwards masking paradigm in which the duration between critical stimulus presentation (ex. a face) and a neutral mask stimulus (ex. a

checkerboard), or SOA, was linearly manipulated around threshold. Behavioral responses mirrored those found in Sergent and Dehaene (2004), showing stimuli were almost never perceived at short SOAs, perceived half of the time at threshold, and perceived almost all of the time at long SOAs. Crucially, perception was not significantly different between the shortest adjacent SOAs (i.e. 16ms and 33ms) and between the longest adjacent SOAs (i.e. 66ms, 83ms, and 100ms), as measured by responses on an objective stimulus discrimination task. Significant differences in objective performance between the 33ms SOA and threshold SOA (50ms), and between threshold SOA and the 66ms SOA demonstrated that perception changes most dramatically around threshold, suggesting that consciousness is likely not modulated directly by stimulus strength. Event related potentials (ERPs) supported this finding, with the P3b exhibiting a non-linear increase in amplitude around threshold. Additional signals – particularly the N2 (discussed further in section 1.4.1) – showed some evidence of a non-linear amplitude increase, though the effect was weaker compared to the P3b due to trial-by-trial participant reports.

Recently, Sergent et al. (2021) implemented a no-report auditory detection paradigm to determine if the non-linear amplitude changes observed in Del Cul et al. (2007) persist for other modalities and in the absence of report. During the "active" listening sessions, participants were instructed to detect the French vowels /a/ and /ə/ embedded in continuous noise. The intensity of the vowel stimuli was increased linearly around perceptual threshold (-9dB), and perception was measured using both an objective 2-alternative-forced-choice (2AFC) vowel discrimination task and a subjective audibility rating task. Both behavioral and neural responses during the active session exhibited bifurcation dynamics, again centering around the P3. As expected, the P3 was absent during "passive" listening sessions in which participants listened to the French vowel

stimuli while completing unrelated tasks such as a visual reaction time task, mind-wandering task, and a mental arithmetic task. Instead, ERP analysis showed a bifurcated late central negativity, which Sergent et al. (2021) argue supports bifurcation dynamics as a signature of the transition from non-conscious to conscious processing.

The findings from Del Cul et al. (2007) and Sergent et al. (2021) strongly suggest that brain signals demonstrating bifurcation dynamics in the absence of behavioral report are likely candidate NCCs. Signals showing non-bifurcated patterns (i.e. linear or non-linear unimodal) or bifurcated patterns only present in report conditions (i.e. the P3b), on the other hand, may be ruled out as potential NCCs (Cohen et al., 2024; Del Cul et al., 2007; Sergent et al., 2021). More evidence is still needed to confirm that bifurcation dynamics are a signature of conscious processing. As noted by Sergent et al. (2021), bifurcation dynamics must be generalizable across modalities, appear for both complex and simple stimuli, and display bimodal (rather than linear or non-linear unimodal) dynamics to be considered a generic signature of conscious perception.

1.3.1 Limitations of Backwards Masking in Investigating Bifurcation Dynamics

Thus far, only Cohen et al. (2024) has combined no-report methodology with linear/non-linear manipulations of stimulus strength and perceptual awareness to identify candidate neural correlates of visual perception. Face stimuli were masked at linearly spaced SOAs (16.67ms, 33.33ms, 50ms, 66.67ms, and 83.33ms) while continuous EEG was recorded and participants completed an objective detection task in the report condition and a probe detection task in the no-report condition. Behavioral reports followed a bifurcated pattern of

perception, reflecting previous findings from Del Cul et al. (2007) and Sergent et al. (2021). Out of several ERP components of interest, two signals revealed bifurcation dynamics: the P3b and a fronto-central negativity beginning around 250ms, termed the fronto-central N2. Consistent with Del Cul et al. (2007), the P3b displayed non-linear bifurcation dynamics, with no significant difference in amplitude between the shortest and longest SOAs, and significant differences in amplitude surrounding the middle SOA. Once again, a bifurcated P3b was absent in the no-report condition, replaced instead by the fronto-central N2. Bifurcation dynamics were not identified for other components of interest, namely the P1 and N170/VAN.

Cohen et al. (2024)'s analysis of the N170/VAN was likely hindered, however, due to the limitations imposed by backwards masking. While a robust and effective method for manipulating awareness of visual stimuli, backwards masking maintains significant drawbacks for isolating signals associated with perceiving target stimuli, particularly in regard to stimulus timing. ERP responses to visual stimuli typically last ~500ms, yet, in most backwards masking studies, mask stimuli are presented 17-100ms after the target stimuli. This results in a significant temporal and spatial overlap between brain responses elicited by the target and mask stimuli, requiring neural responses elicited by the mask to be subtracted from the overall data. To isolate signals elicited solely by target stimuli, ERPs from trials in which only a mask was presented are shifted in time to match the SOA of target-mask trials (i.e. mask-only trials are shifted forward in time by 67ms to match target-mask trials with an SOA of 67ms). The time-shifted mask-only trials are then subtracted from the target-mask trials, ideally yielding target-specific ERPs (Cohen et al., 2024; Del Cul et al., 2007). While various methods have been devised to improve mask-subtraction, lingering effects of the masks often remain present in the subtracted data, particularly

latency differences in ERPs between ~150-300. These latency shifts complicate measuring earlier components like the N1, VAN, and N170 (Cohen et al., 2024; Cohen et al., 2020). Additionally, seen-versus-unseen comparisons post-mask subtraction may reflect not only perceptual differences, but residual differences in brain regions like the prefrontal cortex (PFC) that manage perceptual conflict and competition, cognitive demands, etc. More simply, on "seen" trials participants consciously perceive the stimulus quickly followed by the mask (two stimuli competing for cognitive processing resources) while on "unseen" trials, they only consciously perceive the mask (one stimulus, no competition for cognitive resources). Cohen et al. (2024)'s use of faces as target stimuli also add increased complexity to the design, potentially leading to the presence of signals related to face perception specifically rather than more basic conscious awareness (e.g., of simple visual patterns). By simplifying the stimulus and perceptual manipulation used by Cohen et al. (2024), future studies may confirm the fronto-central N2 as a candidate NCC and potentially identify additional candidate NCCs.

1.4 Candidate NCCs

Several candidate NCCs have been identified in the literature. For the purposes of this study, only event related potentials (ERPs) and metastability are discussed.

1.4.1 Event Related Potentials

Visual Awareness Negativity

The visual awareness negativity (VAN) has been identified amongst several studies as a potential NCC. The VAN is relatively more negative in response to seen stimuli (when compared to unseen stimuli), and is typically observed from ~100-300ms in occipital electrodes (100ms represents the earliest onset of the VAN, with peak amplitudes typically around 200-250ms). The intensity of the VAN has also been shown to be maximal at sites contralateral to the presented stimulus (Dembski et al., 2021; Förster et al., 2020). The VAN has been associated with processing in early visual areas and implicated in sensory theories of consciousness relating recurrent processing to consciousness (Koch et al., 2016; Förster et al., 2020). Evidence supporting the VAN as a candidate NCC is relatively robust, with few studies failing to observe the component under expected conditions. Upon critical analysis, these studies neglected to account for the lateralization of VAN topography in response to stimuli presented in one hemifield, averaging data across both hemifields and functionally negating any significant effects (Dembski et al., 2021). Depending on the specific manipulations or paradigms used, several additional components can occur during the time window of the VAN, such as the N170 (specific to face stimuli) and the N2pc (specific to ipsilateral vs. contralateral control of spatial attention), though it is still unclear whether these signals are variations on the VAN or separate signals all together.

Fronto-Central N2

The fronto-central N2 was recently identified by Cohen et al. (2024) as a potential candidate NCC. Several additional studies utilizing a variety of methods have identified signals with similar spatial and temporal characteristics to the fronto-central N2 (Cogitate et al., 2023; Cohen et al., 2020; Dellert et al., 2022; Gaillard et al., 2009; Schlossmacher et al., 2020; Sergent et al., 2021). The fronto-central N2 is typically observed in central anterior electrodes from ~250-300ms, and has maximal amplitude in response to seen stimuli. Sergent et al. (2021) observed a burst in inter-trial EEG variability for threshold level stimuli during the approximate time window of the fronto-central N2, further supporting its potential as a candidate NCC. While the time-windows of the VAN and fronto-central N2 partially overlap, differences in scalp topography and specific modality response distinguish the two components (Förster et al., 2020; Cohen et al., 2024). The fronto-central N2 has been proposed to index a stage of processing between early sensory activation and global information sharing with cognitive systems (for further discussion on early recurrent processing and global "ignition", see section 1.5), possibly representing a signature of conscious access.

1.4.2 Metastability

Understanding how the brain differentially encodes and processes content may also aid in identifying neural correlates of awareness. Researchers have previously relied on indirect behavioral measures, such as differences in subjects' reaction times to varying stimulus and task conditions (known as 'mental chronometry'), to map the temporal dynamics of cognitive processing stages (Pashler, 1994; Sternberg, 1969). Indirect behavioral measures, while informative,

are unable to resolve whether processes are serial or parallel, or if processing occurs in discrete versus continuous stages (Dehaene & King, 2016). Direct brain-imaging techniques like EEG, MEG, and fMRI have since enabled researchers to resolve the ambiguities present in behavioral measures of cognitive processing, especially when paired with robust machine learning techniques capable of handling large amounts of complex neural data (King & Dehaene 2014; Dehaene & King, 2016). By applying machine learning algorithms to brain-imaging data, researchers can predict, or "decode," perceptual content and piece together temporal dynamics of cognitive processes (King & Dehaene 2014; Dehaene & King, 2016; Grootswagers et al., 2017).

Multivariate Pattern Analysis

Decoding, or multivariate pattern analysis (MVPA), was first applied to fMRI data as a more comprehensive alternative to univariate analyses (Haxby et al., 2001). When applied to time series data (i.e. EEG and MEG, as opposed to spatial data from fMRI), MVPA can reveal when mental representations are activated (King & Dehaene, 2014). Unlike univariate analyses, MVPA takes multiple variables into account (i.e. electrodes/sensors in EEG and MEG or voxels in fMRI) rather than the relative strength between conditions at each individual point of measurement. Multivariate decoding methods are also advantageous, as they increase signal-to-noise ratios and reduce spatial variability across subjects, all while preventing "double-dipping" of data (i.e. separating inferential data from confirmatory data) (Dehaene & King, 2016).

The standard approach is to train machine learning pattern classifiers to extract information about specific characteristics of a stimulus or task (e.g., stimulus presence, stimulus orientation, task relevance, etc.) from time slices

(e.g., time *t*) of electrophysiological data. The trained pattern classifiers are then tested on a separate slice of data at the same time points (time *t*), with the resulting accuracy of the pattern classifiers plotted over time. Decoding accuracy of the pattern classifiers typically begins at chance (50% accuracy) before stimulus onset, spikes shortly after stimulus onset, then gradually decays back to chance (Figure 2). Because pattern classifiers are trained to distinguish between stimulus conditions (e.g., classifiers are trained to distinguish brain responses to a green square from responses a blue circle, then tested to make the same distinction based on new, unlabeled data), above chance decoding performance indicates the brain differentially encodes information about a given stimulus. Visualizing the decoding accuracy of the trained pattern classifiers, then, allows researchers to track when different aspects of the stimuli or task are encoded in brain activity over the course of a given trial. One limitation of MVPA is the inability to know if the differential brain activity picked up by the classifiers are used directly by the brain, or if they are the result of some secondary or indirect process. By applying linear classifiers like a Support Vector Machine (SVM), the likelihood that only explicit neural codes are decoded increases (Dehaene & King, 2016).

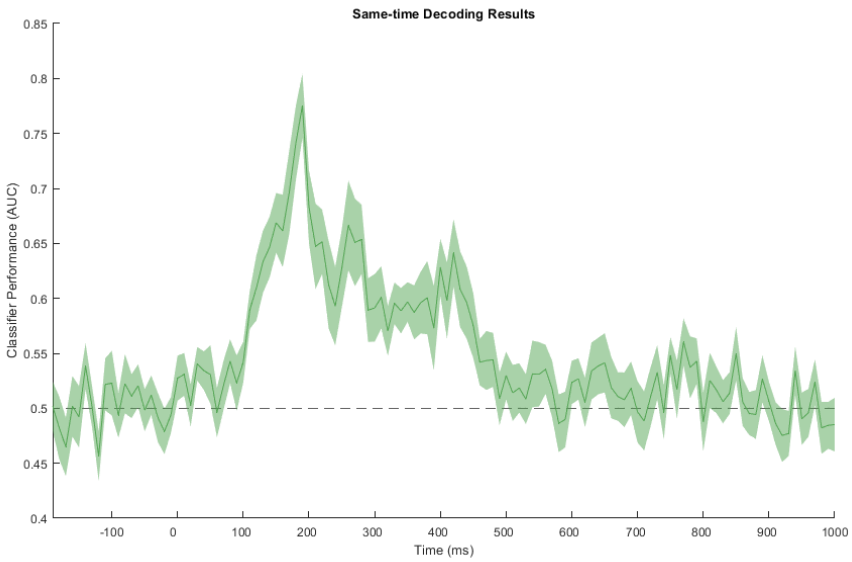

Figure 2: Example of same-time decoding of face stimulus visibility in an inattentional blindness paradigm.
Classifier performance was plotted as a function of the area under the receiver operating curve (AUC), with chance performance (50%) shown by the dotted line. A spike in classifier performance beginning at ~150ms suggests a difference in brain activity between perceptually visible and invisible faces.

Thus far, studies have utilized MVPA with EEG to decode a variety of stimulus features, including location, orientation, presence, visibility, and size (King et al., 2016). Higher level perceptual and cognitive features can also be decoded, such as object category and task (King and Dehaene, 2014). Decoding stimulus presence is particularly useful for identifying potential NCCs. By training classifiers to distinguish between stimulus present and stimulus absent trials at each stimulus visibility level (e.g., at each variable SOA or contrast level), for instance, the onset of conscious processing can be determined (Cohen et al., 2024). In this way, multivariate decoding directly complements many ERP methods, providing fine grain differences in brain activity in response to perceived versus not perceived stimuli.

Temporal Generalization of Decoding

Multivariate decoding methods can be extended even further to investigate how neural codes generalize or recur over time. Using the temporal generalization method, classifiers are trained on a subset of data at time *t*, and then tested on all other time points, or time *t'*. The resulting temporal generalization matrix plots training time in rows (y-axis) and testing time in columns (x-axis), with decoding accuracy visualized via color mapping. Decoding accuracy along the diagonal of the matrix is analogous to the same-time decoding curve described above, with the classifier trained and tested on the same data point (known as "diagonal decoding"). Patterns of above chance decoding within temporal generalization matrices can hint at underlying brain processes (King & Dehaene, 2014; Dehaene & King, 2016) (Figure 3). Decoding off the diagonal, for example, reflects consistent brain patterns across time points, as demonstrated by the classifiers' ability to use information earlier in time to decode information later in time (and vice versa). Additionally, cross-condition decoding may reveal how brain processing changes as perceptual content or task changes (King & Dehaene, 2014).

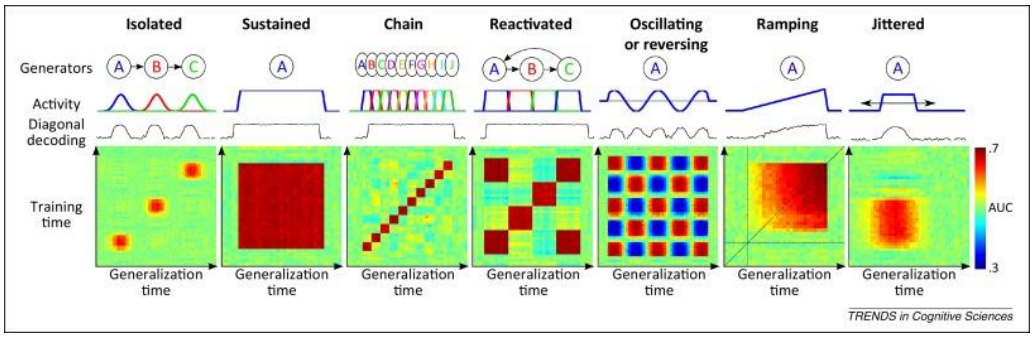

Figure 3: Example temporal generalization matrices derived from differing neural generators (King & Dehaene, 2014).

Several neural generators were simulated to determine the resulting temporal generalization matrices. King, J.-R., & Dehaene, S. (2014). Characterizing the dynamics of mental representations: The temporal generalization method. Trends in Cognitive Sciences, 18(4), 203–210. https://doi.org/10.1016/j.tics.2014.01.002.

A number of studies have identified characteristic temporal generalization patterns associated with awareness. Dehaene & King (2016) suggest late metastable activity as a signature of consciousness, characterized by a late, square-shaped pattern of above chance decoding. This late square-shaped pattern is thought to represent the transition of amplified information into a metastable representation possibly associated with conscious access (Dehaene & Changeux, 2011; Kouider et al., 2013; Salti et al., 2015; Schurger et al., 2015). Metastable patterns of decoding are typically preceded by extended diagonal decoding, although this relationship is largely mediated by perceptual awareness. When applied to variably masked stimuli (Charles et al., 2014; Del Cul et al., 2007), temporal generalization reveals early diagonal decoding (~150-350ms) for subliminal stimuli, which temporally extends for threshold stimuli and finally spreads into a late square-shaped pattern for above-threshold stimuli (~350-600ms). Similar results were found when linearly degrading auditory stimuli (Sergent et al., 2021). Crucially, these findings only extend to report

paradigms, with several studies largely failing to identify late square-shaped patterns in no-report paradigms (Cohen et al., 2024; Sergent et al., 2021; Zhu et al., 2024).

1.5 Sensory Versus Cognitive Theories of Consciousness

While the search for NCCs is largely theory-neutral, findings from NCC focused studies maintain major theoretical implications. Theories of consciousness are particularly helpful in situating the topography, timing, and specificities of NCCs, providing a framework for understanding *how* and *when* consciousness occurs. As previously discussed, the need to isolate pre-conscious processes and post-perceptual processes from consciousness itself has complicated the search for NCCs. Situating NCC research within a theoretical framework can help alleviate some of these complications, albeit not entirely, and provide explanatory insight (Seth and Bayne, 2022).

Theories of consciousness generally fall into two camps: sensory or cognitive. Both categories of theory seek to close the explanatory gap between neural activity and consciousness (also known as the "Hard Problem" of consciousness – for further discussion, see Chalmers (1995) and Block (2002)). Sensory theories of consciousness, including Recurrent Processing Theory and Integrated Information Theory (IIT), place an emphasis on the role of sensory areas of the brain in enabling awareness (Koch et al., 2016; Boly et al., 2017; Lamme, 2018). These theories focus on early, sensory localized signals (originating in parietal, temporal, and occipital cortices) to support the integration of sensory information or re-entrant processing within perceptual

cortices (Seth and Bayne, 2022; Förster et al., 2020). Cognitive theories of consciousness, including Higher-Order Thought Theory (HOT) and Global Neuronal Workspace Theory (GNWT), place an emphasis on the role of higher-level brain areas in enabling awareness (Dehaene, 2014; Odegaard et al., 2017; Brown et al., 2019; Mashour et al., 2020). These theories focus on late, fronto-parietal networks in allowing sensory information to become the target of meta-representations, or accessed by the global workspace. Signals like the P3b have largely been cited as supporting theories like GNWT, but have more recently been attributed to post-perceptual processes rather than consciousness itself. The PFC may still play a significant role in consciousness, though (Panagiotaropoulus, 2024). GNWT also supports late metastability as a marker of consciousness, citing spread, "square-shaped" patterns of temporal generalization as evidence of sensory information being broadcast to the global workspace (Dehaene and King, 2016).

As previously noted, the fronto-central N2 sits at a cross-roads between sensory and cognitive theories of consciousness. The relatively late timing and anterior topography of the fronto-central N2 contrast with predictions made by sensory theories of consciousness, which are often associated with early electrophysiological signals like the VAN (Förster et al., 2020; Dembski et al., 2021). Similarly, its bilateral posterior positive foci and short time window contrast with predictions made by cognitive theories of consciousness, which are often associated with sustained response time and broad spatial distribution. Such evidence suggests that the fronto-central N2 represents a stage of perceptual processing between sensory representations and cognitive processing of those representations.

Taken together, consciousness may occur as a sequence of processes encapsulating both sensory and cognitive theories of consciousness. Under this

view, sensory focused signals like the VAN may represent the "packaging" of sensory information into an accessible form, with the fronto-central N2 indexing the access of that information by higher order processes (e.g, for decision making, judgement, internal thought, etc.). Consciousness, then, serves as a way to make organisms poised for action, allowing for flexibility in how sensory information is processed and utilized. Further evidence is undeniably needed to assess these claims, which the current study seeks to accomplish.

1.6 The Current Study

The current study seeks to investigate potential NCCs through the use of bifurcation dynamics. We aim to replicate the findings of Cohen et al. (2024) using a simplified design to prevent complications brought on by backwards masking. To do so, the visibility of simple sinewave "Gabor" gratings embedded in dynamic noise will be linearly degraded while continuous EEG is recorded. To isolate awareness from post-perceptual processes, we will use a no-report design paired with a behavioral report condition to confirm participants' perception of the Gabor gratings was bifurcated.

In line with Cohen et al. (2024), we predict a non-linear bifurcated signal in central anterior electrodes (i.e. the fronto-central N2) in only the no-report condition. Due to latency differences related to mask subtraction, Cohen et al. (2024) were unable to fully analyze potential bifurcation dynamics of the visual awareness negativity (VAN), so our analysis of the VAN will be completely novel. In the report condition, we predict non-linear amplitude increases in the P3b signal, reflecting late-stage bifurcation around perceptual threshold. In the no-report condition, we predict the P3b will be absent.

Additionally, temporal generalization of decoding will show a late metastable representation associated with conscious perception in the report condition, with some off-diagonal decoding associated with awareness in the no-report condition. We also predict accuracy of stimulus presence decoding in the no-report condition will increase as a function of visibility. Following behavioral bifurcation dynamics, above chance decoding accuracy will only be observed when classifying between threshold level stimuli and adjacent stimuli (i.e. stimuli just below or just above perceptual threshold), reflecting no difference in conscious processing between minimally or maximally visible stimuli. If possible, decoding of stimulus orientation will follow a similar pattern to stimulus presence decoding, with an increase in accuracy as visibility increases.

Chapter 2: Methods

2.1 Participants

Fifty five total subjects were recruited for participation in the study. 6 subjects were recruited for participation in the first pilot study, with 1 subject excluded due to a false alarm rate of 27% (n=5). After the first pilot study, 6 additional subjects were recruited for participation in the second pilot study, with 1 subject excluded due to an unexpected error in the experiment code (n=5). 43 subjects were recruited for participation in the main experiment. 13 failed to meet exclusion criteria, resulting in a total sample of 30 participants in the main experiment. Participants were between 18 to 35 years old, with no pre-existing neurological conditions (i.e. epilepsy or other seizure disorders), no concussions in the past year, and no lazy eye. Participants were compensated $30 for their time in the main experiment, and $50 for participation in the second pilot study due to the task being split into two EEG sessions.

2.2 Stimuli

Simple sine wave "Gabor" gratings embedded in continuous dynamic noise (similar to TV static), were used to simplify the design implemented by Cohen et al. (2024). The Gabor gratings were two cycles per degree of visual angle, and oriented either 45° ("right" of vertical) or 135° ("left" of vertical) (Figure 4). Additionally, the Gabor gratings appeared with an inter-trial temporal jitter of 400ms and spatial jitter of 1 degree of visual angle relative to

their center. The dynamic noise was presented at 12 degrees of visual angle, with a spatial frequency of 2 pixels per degree and a refresh rate of 0.033 seconds (Figure 5). All stimuli were presented at 120Hz on a 1920 x 1080 px BenQ monitor, positioned at eye level approximately 57 cm in front of the seated participant. Stimuli were be rendered and presented using Psychtoolbox 3.0.14 (Kleiner et al., 2007) for MATLAB (v. 2022b).

Figure 4: Example of a "left" Gabor grating (135˚).
The Gabor grating shown above is set to a contrast level of -0.8 (arbitrary units) for demonstration purposes.

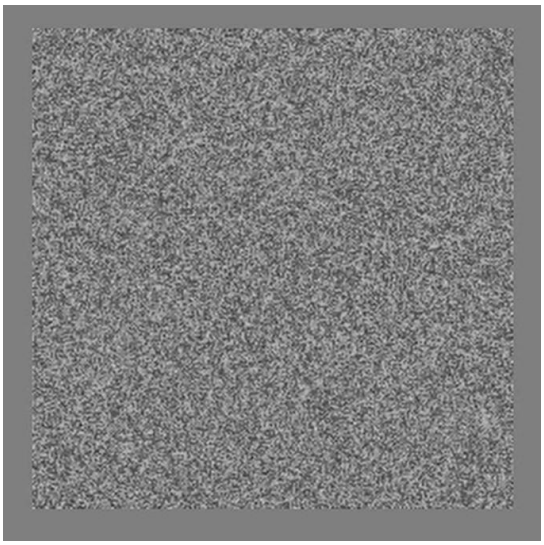

Figure 5: Still example of continuous dynamic noise.
Gabor gratings were embedded in the continuous dynamic noise shown above to aid in gradually degrading their visibility. The dynamic noise was presented at 12 degrees of visual angle, with a spatial frequency of 2 pixels per degree and a refresh rate of 0.033 seconds.

2.3 Design

2.3.1 Thresholding Procedure

Perceptual thresholds were established using varied contrast levels to degrade the visibility of the Gabor gratings. Threshold contrast values were tailored to each participant using the QUEST staircase procedure (Watson et al., 1983), consisting of 3 blocks of 50 trials of a 2-alternative-forced-choice (2AFC) task orientation discrimination task in which participants indicated the direction of the Gabor (i.e. left-tilting versus right-tilting). Contrast values obtained after the final two QUEST blocks were averaged to determine the final threshold contrast value. The remaining contrast levels were linearly spaced out from the threshold contrast value (contrast level 3) using a step size of 0.12 (arbitrary

units). Linearly stepping out from threshold resulted in a bifurcated behavioral response resembling a standard sigmoid-shaped psychometric curve, with below chance accuracy on the 2AFC discrimination task for Gabor gratings below threshold, at chance accuracy for Gabor gratings at threshold, and above chance accuracy for Gabor gratings above threshold. Despite their linear spacing, there were minimal differences in visibility between the first two and last two contrast levels, and maximal differences in visibility for contrast levels adjacent to threshold.

2.3.2 Pilot Experiment 1

Initial pilot testing was performed to confirm that the behavioral visibility curves showed the expected sigmoid shape, that there were sufficient numbers of trials to obtain sufficient signal-to-noise ratios in the ERP data, and that the P3b was present in the report condition and absent in the no-report condition. The first pilot experiment (n=5) was broken into two conditions: report and no-report. The report condition consisted of 20 blocks of 30 trials, while the no-report condition consisted of 24 blocks of 30 trials, resulting in 100 trials per contrast level (including blank trials, or trials during which no Gabor is presented). The number of blocks were increased in the no-report condition to account for probe trials, which were presented on a subset of trials (~20%). Participants were instructed to attend to the dynamic noise and detect a coherent motion probe (60% coherence) presented for 300ms within the same spatial extent as the Gabor gratings to ensure that the gratings remained task-irrelevant while maintaining subjects' attention. In the report condition, participants completed a 2AFC orientation discrimination ("left" arrow key or "right" arrow key) and perceptual awareness scale (PAS, 1-4) rating task, allowing for objective and subjective

measures of perception to be collected (Ramsøy & Overgaard, 2004). The report and no-report conditions were balanced across participants so that half of the participants completed the report condition first (n=3) and half of the participants will complete the no-report condition first (n=3).

Participants were familiarized with the stimuli prior to the thresholding procedure (see Thresholding Procedure above), which always preceded the report or no-report tasks. To achieve 75% accuracy on the orientation discrimination task, QUEST was set to threshold to 70% accuracy, compensating for possible learning over the course of the experiment. Four additional contrast levels were established based on the threshold contrast value, linearly stepped out by 0.12 (arbitrary units) to yield five total contrast levels (contrast level 1, level 2 below threshold, level 3 at threshold, and level 4 and level 5 above threshold). Brain activity was recorded with EEG at 1000Hz using a custom 64-channel electrode cap during both report and no-report conditions. One participant was excluded from pilot analysis due to a false alarm rate higher than 20%, as determined by PAS rating higher than 1 (i.e., reporting seeing stimuli on blank trials in which no stimulus was physically presented).

Average behavioral data from the first pilot experiment demonstrated a promising bifurcated pattern reflected in both objective (Figure 6) and subjective measures of awareness (Figure 7). Despite the overall bifurcated pattern, threshold contrast level accuracy on the 2AFC task was higher than desired. This elevated threshold accuracy was likely due to learning that occurred between QUEST and the report condition.

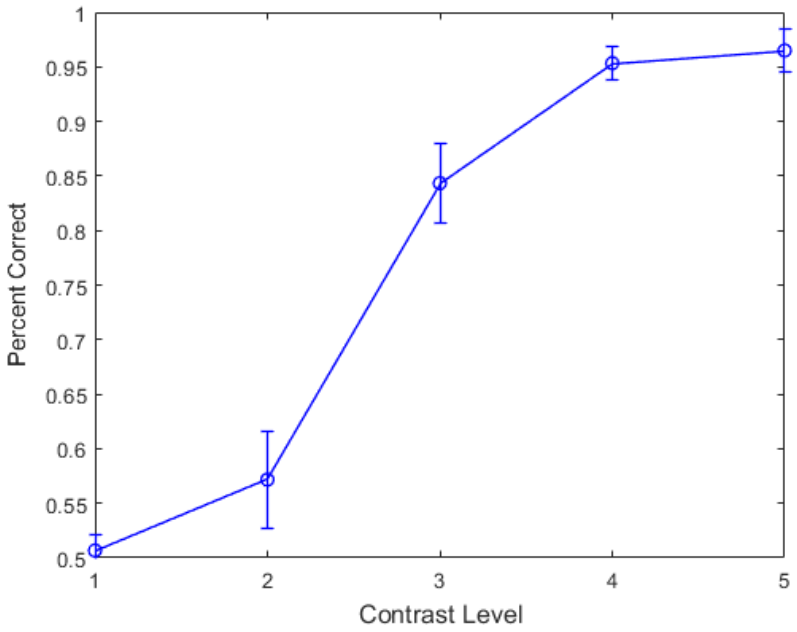

Figure 6: Average orientation discrimination task performance from the first pilot experiment.

Results from the 2AFC orientation discrimination task from the first pilot experiment plotted by contrast level and averaged across participants. Threshold (contrast level 3) accuracy was 84.3%, higher than the desired 75% after QUEST thresholding. Error bars represent standard error of the mean.

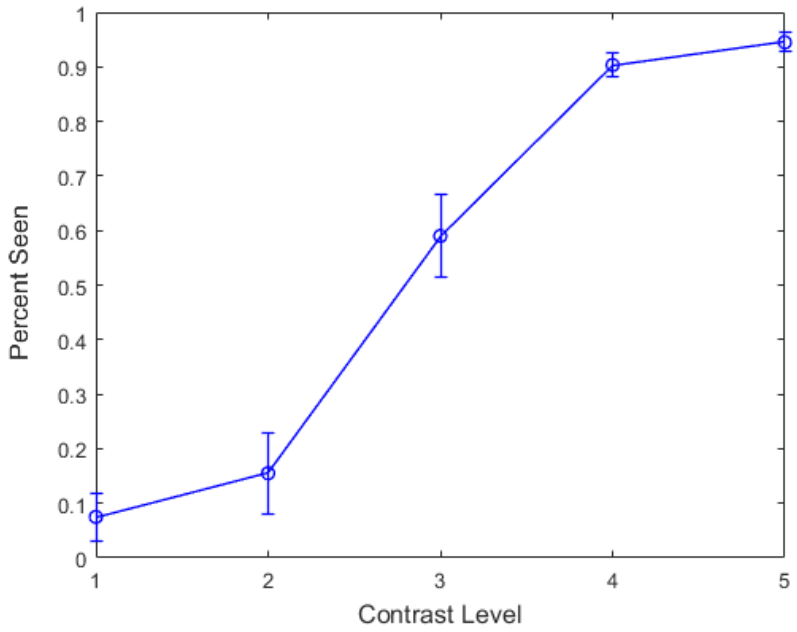

Figure 7: Average percentage of Gabor gratings rated as seen during the first pilot experiment.

PAS ratings of 2 or higher indicated a Gabor grating was seen. Data were averaged across participants to yield the above behavioral function. Error bars represent standard error of the mean.

Preliminary ERPs from the first pilot experiment were also examined. As expected, scalp distributions for each participant showed a consistent P3b in the report condition. Data from the no-report condition was less clear, showing a much smaller late positivity compared to the report condition that was distributed over similar electrodes as the P3b, albeit somewhat more posterior (at least for several subjects) (Figure 8). Additional signals of interest were present in the preliminary ERPs, namely the VAN. A distinct VAN was observed for contrast level 4 and 5 stimuli, yet largely attenuated for threshold level stimuli (about -1mV). As expected, no VAN was observed for contrast level 1 and 2 stimuli.

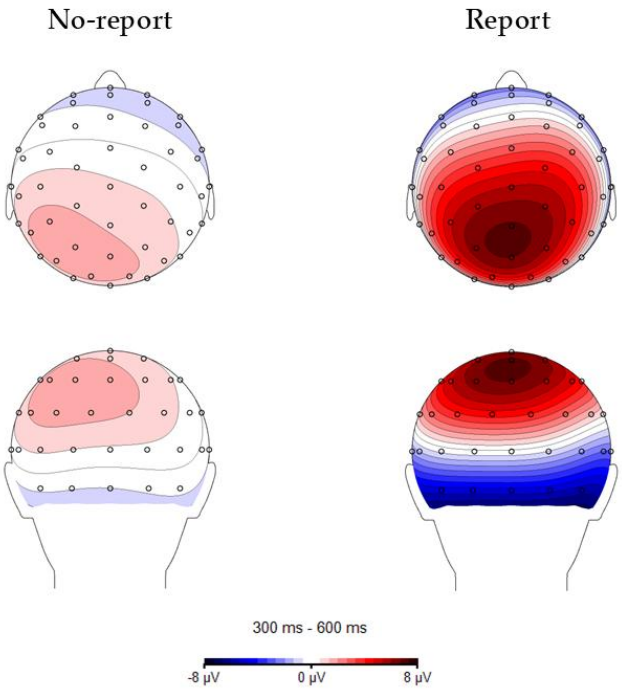

Figure 8: Topographic scalp distributions from the no-report and report conditions of the first pilot experiment.
EEG data were preprocessed to examine preliminary ERPs from the first pilot experiment. Scalp maps show evoked potentials from 300-600ms (typical P3b time-window) for contrast level 5 using an average reference.

Changes to the experimental design were made based on results from the first pilot experiment and implemented in a second pilot experiment.

2.3.3 Pilot Experiment 2

The second pilot experiment (n=5) utilized the same basic design as the first pilot experiment. Participants first completed the QUEST procedure, then completed either the report condition (n=3) or the no-report condition (n=2) first. To adjust for the elevated threshold accuracy observed during the first pilot experiment, QUEST was set to threshold to 65%. Five additional contrast levels

were established based on the threshold contrast value, linearly stepped out by 0.12 (arbitrary units) to yield six total contrast levels (contrast level 1, level 2 below threshold, level 3 at threshold, and level 4, level 5, and level 6 above threshold). By adding an additional contrast level above threshold, we sought to address possible attention-related changes in stimulus sensitivity between the report and no-report conditions. The report condition consisted of 17 blocks of 42 trials, while the no-report condition consisted of 20 blocks of 42 trials, resulting in 102 trials per contrast level (including blank trials). The number of blocks were again increased in the no-report condition to account for probe trials, which were presented on a subset of trials (~17%). Participants were instructed to attend to the dynamic noise and detect a red dot probe with a radius of 12 pixels (RGB: 255, 0, 0) presented for 100ms in a randomized location within the same spatial extent as the Gabor gratings. A red dot probe was used instead of the coherent motion probe to ensure that the Gabor gratings remained task-irrelevant, as attending to coherent motion in the dynamic noise may have inadvertently rendered the gratings partially task-relevant during the first pilot experiment.

The second pilot experiment was split into two EEG recording sessions, one for the report condition and one for the no-report condition, in order to shorten the amount of time each participant was actively engaged in a task. Brain activity was recorded with EEG at 1000Hz using a custom 64-channel electrode cap during both sessions.

Behavioral data from the second pilot experiment were consistent with those of the first pilot experiment, showing a promising pattern of bifurcation for both objective (Figure 9) and subjective measures of awareness (Figure 10). 2AFC results showed a decrease in threshold contrast level accuracy, likely due to the reduction in QUEST thresholding implemented after pilot experiment 1.

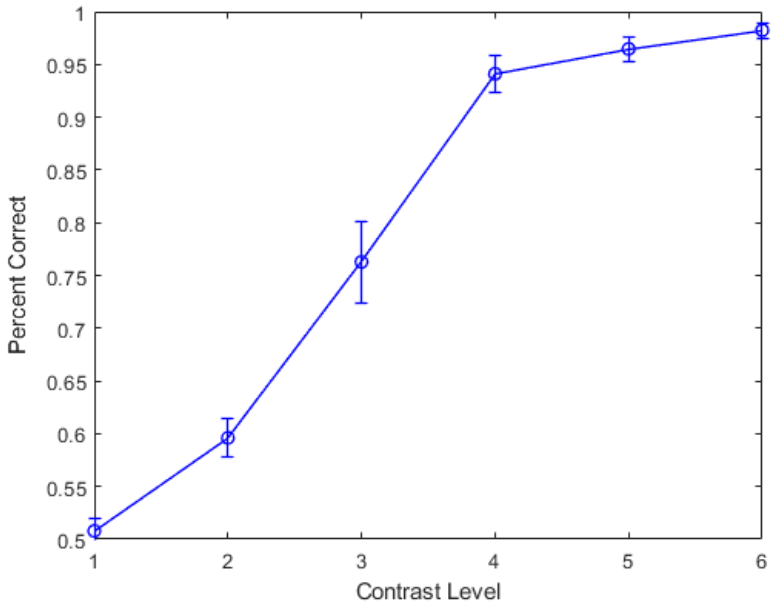

Figure 9: Average orientation discrimination task performance from the second pilot experiment.

Results from the 2AFC orientation discrimination task from the second pilot experiment plotted by contrast level and averaged across participants. Error bars represent standard error of the mean. Average threshold accuracy was 76%, demonstrating that changes to the QUEST thresholding worked as expected.

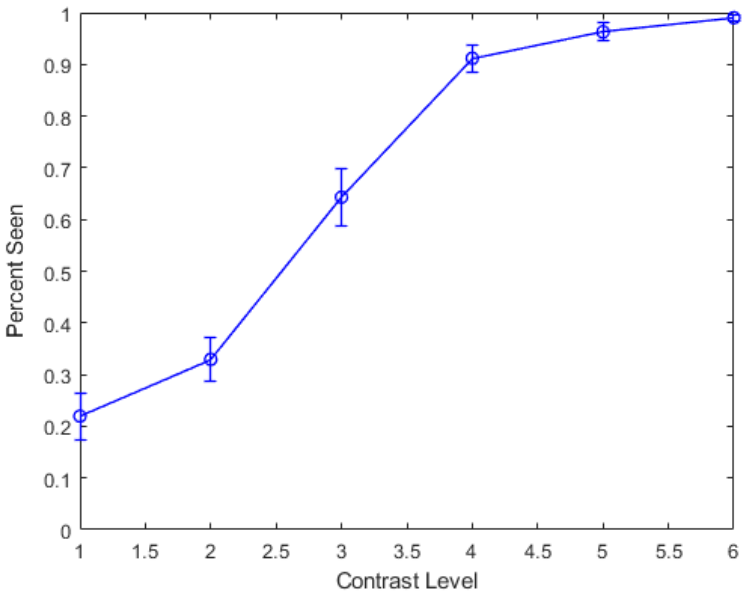

Figure 10: Average percentage of Gabor gratings rated as seen during the second pilot experiment.

PAS ratings of 2 or higher indicated a Gabor grating was seen. Data were averaged across participants to yield the above behavioral function. Error bars represent standard error of the mean.

To determine if there was a differential effect of the red dot probe compared to the coherent motion probe, preliminary ERPs from the second experiment were examined (Figure 11). Scalp topography consistent with the P3b was observed for the report condition, as was shown in the first pilot experiment. Scalp topography from the no-report condition was also consistent with the first pilot experiment, showing a late positivity located slightly more posterior compared to a typical P3, albeit during the same time window. No major difference was observed in VAN amplitude at threshold, with evoked potentials from contrast level 6 maintaining a similar amplitude to those from contrast levels 4 and 5.

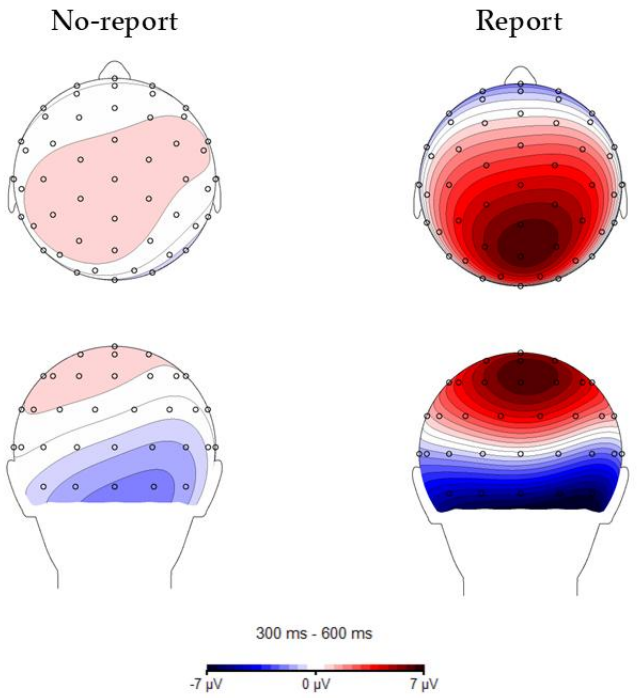

Figure 11: Topographic scalp distributions from the no-report and report conditions of the second pilot experiment.
EEG data were preprocessed to examine preliminary ERPs from the second pilot experiment. Scalp maps show evoked potentials from 300-600ms (typical P3b time-window) for contrast level 5 using an average reference.

2.3.4 Main Experiment

Parameters for the main experimental design (Figure 12) were determined from the results of the first two pilot experiments. Participants first completed the QUEST thresholding procedure, which was set to threshold to 70% accuracy based on behavioral results from the second pilot experiment. Four additional contrast levels were established based on the threshold contrast value, linearly stepped out by 0.12 (arbitrary units) to yield five total contrast levels (contrast level 1, level 2 below threshold, level 3 at threshold, and level 4 and level 5 above

threshold). Participants then completed the main no-report condition, which consisted of 26 blocks of 36 trials, resulting in 130 trials per contrast level (including blank trials). During the no-report condition, participants were instructed to attend to the dynamic noise and detect infrequent red dot probes (identical to those used in the second pilot experiment) which appeared on a subset of trials (~16%). Continuous 64-channel EEG was recorded during the no-report condition, sampled at 1000Hz. After completing the main no-report condition, the EEG recording was stopped, and participants completed an abbreviated report condition to assess awareness of the Gabor gratings and obtain behavioral psychometric functions. The behavioral report condition consisted of 8 blocks of 30 trials, resulting in 40 trials per contrast level (including blank trials). In all, the main experiment design improved on the design of the previous two pilot experiments by only recording EEG data in the no-report condition, increasing the number of trials in the no-report condition to improve signal-to-noise ratios, and thresholding to 70% to account for perceptual learning.

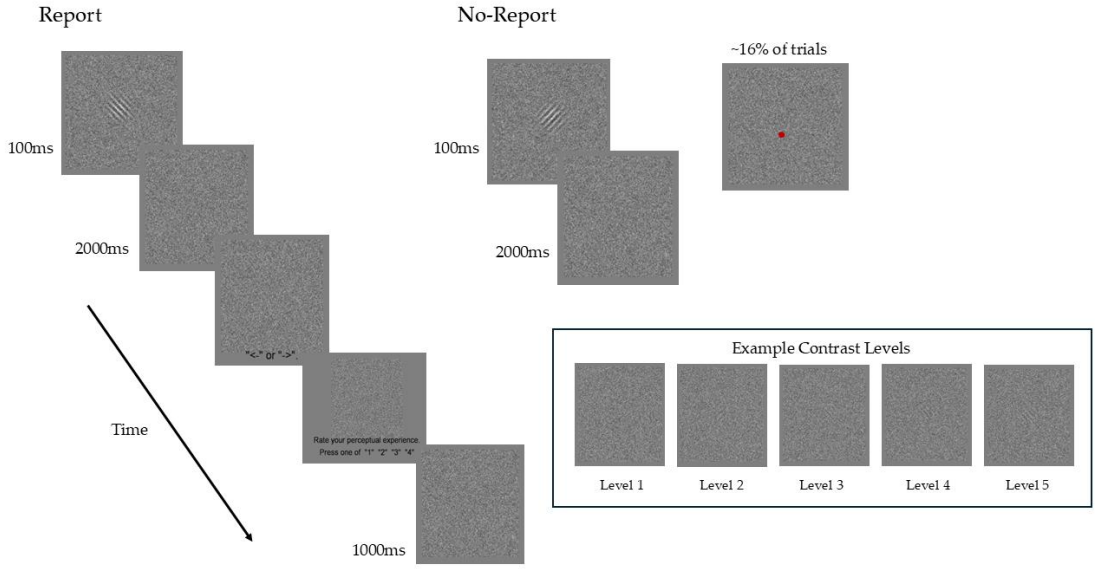

Figure 12: Design of the final stimulus degradation paradigm.

A Gabor grating (or blank) was presented for 100ms embedded in dynamic noise, followed by the dynamic noise only for 2000ms, with a temporal jitter between trials of 400ms and a spatial jitter of 1 degree of visual angle across different trials. In the report condition, subjects were instructed to indicate whether the Gabor was oriented to the left ("left" arrow key) or to the right ("right" arrow key), then rate their perceptual awareness of the stimuli using the PAS scale (1-4). Subjects were instructed to respond with a 1 on the PAS scale if they had no experience of the Gabor gratings, with a 2 if they experienced a brief change in the dynamic noise without any idea as to the direction of the Gabor, with a 3 if their experience of the Gabor was almost clear (i.e. the Gabor grating was fuzzy but the orientation was clear), and a 4 if their experience of the Gabor and its orientation were the most salient and clear. In the no-report condition, subjects attended to the dynamic noise and detected infrequent red dot probes (~16% of trials) by pressing the "up" arrow key. The bottom right panel shows example Gabors at each contrast level set from a threshold contrast value of 0.8 (arbitrary units).

2.4 Analysis

2.4.1 Behavioral

Reponses during the report condition were plotted against contrast level and fit using a non-linear regression to determine if behavioral data display bifurcation dynamics. Both objective and subjective measures of perception were analyzed for bifurcation dynamics. The average distribution of PAS responses for each contrast level were also examined to determine if perception of the Gabor gratings followed a bimodal pattern. D' was also calculated for each participant based on PAS responses, then averaged to determine average stimulus sensitivity for each contrast level.

2.4.2 Data Segmentation and Pre-Processing

Continuous EEG data were segmented and preprocessed for further analysis using BrainVision Analyzer (v.2.2). Data were first down-sampled to 500Hz, then visually inspected for consistently noisy or erratic channels. Any such channels were corrected with topographic interpolation using a spherical spline method. Stimulus trigger codes were then time-corrected (10ms) to match display timings. Data re-referenced to both mastoid channels, and a 60Hz notch filter and 0.1Hz high-pass filter were applied. Blinks and eye movements were corrected using independent component analysis (ICA), with remaining artifacts (i.e. muscle tension, noisy electrodes, etc.) removed using a semi-automatic artifact rejection procedure after segmenting trials into 2.0 second stimulus time-locked segments centered around time zero (stimulus onset). Trials were

separated according to contrast level (along with blank trials), then finally baseline corrected (-200ms to 0ms) and averaged.

2.4.3 Multivariate Pattern Analysis

Decoding

Decoding was performed using the MVPA Light decoding toolbox (Treder, 2020) for MATLAB (v. 2023b). Data were exported from BrainVision Analyzer and properly formatted using EEGLAB (Delorme & Makeig, 2004), then down-sampled to 100Hz and z-scored using MVPA Light's default preprocessing parameters. A linear SVM classifier was trained at each contrast level for both report and no-report tasks to distinguish between stimulus present and stimulus absent trials using a 5-fold cross-validation procedure. The data were split into training and testing sets at each time point and independently subaveraged (4 trials per subaverage), with 80% of trials used for training and the remaining 20% used for testing. This process was repeated 5 times so that the decoder was trained and tested on the entire dataset, but classifier performance never assessed on data used for training. Trial numbers were balanced using MVPA Light's "undersampling" parameter so that there were equal numbers of stimulus present and stimulus absent trials in each fold. Classifier sensitivity was evaluated using the area under the receiver operating curve (AUC), and was averaged across folds to quantify overall performance. Additional comparisons used the same methods described above to assess classifier performance on discriminating between adjacent contrast levels (i.e. decoding between contrast level 1 and 2, between contrast level 2 and 3, etc.), between seen and unseen conditions (i.e. contrast level 5 vs. level 1 and contrast level 4 vs. level 2), and

between stimulus orientation (i.e. right oriented Gabors or left oriented Gabors) separately at each contrast level.

Temporal Generalization of Decoding

The temporal generalization of decoding technique was used to assess the neural dynamics of perceptual representations over time. A linear SVM classifier was trained on data from one time point (i.e. 300ms) and then tested on discriminating between classes at all remaining time points, generating a temporal generalization matrix. Classification sensitivity was statistically tested against chance (50%) and corrected for multiple comparisons using 10,000-permutation cluster-based testing. All comparisons made during multivariate pattern analysis were translated to the temporal generalization technique.

Chapter 3: Results

3.1 Behavioral Results

Participant responses on the 2AFC orientation task and PAS rating task during the behavioral report condition were analyzed (Figure 13). A dependent samples t-test was performed to compare responses between adjacent contrast levels (i.e. between levels 1 and 2, 2 and 3, etc.) for each behavioral measure. Mean 2AFC accuracy across subjects (N=30) showed no significant difference between contrast levels 1 (M=0.5188, SD=0.0704) and 2 (M=0.5528, SD=0.0704, $t(29)=-1.6773$, $p=0.1042$), a significant difference between contrast levels 2 and 3 (M=0.7542, SD=0.0807, $t(29)=-13.1988$, $p<0.0005$), a significant difference between contrast levels 3 and 4 (M=0.9208, SD=0.0543, $t(29)=-11.5971$, $p<0.0005$), and a slight significant difference between contrast levels 4 and 5 (M=0.9472, SD=0.0457, $t(29)=-3.1368$, $p=0.0039$). The lack of a qualitative difference in 2AFC accuracy between contrast level 4 and level 5 provided evidence for perceptual bifurcation, despite there being a small quantitatively significant difference.

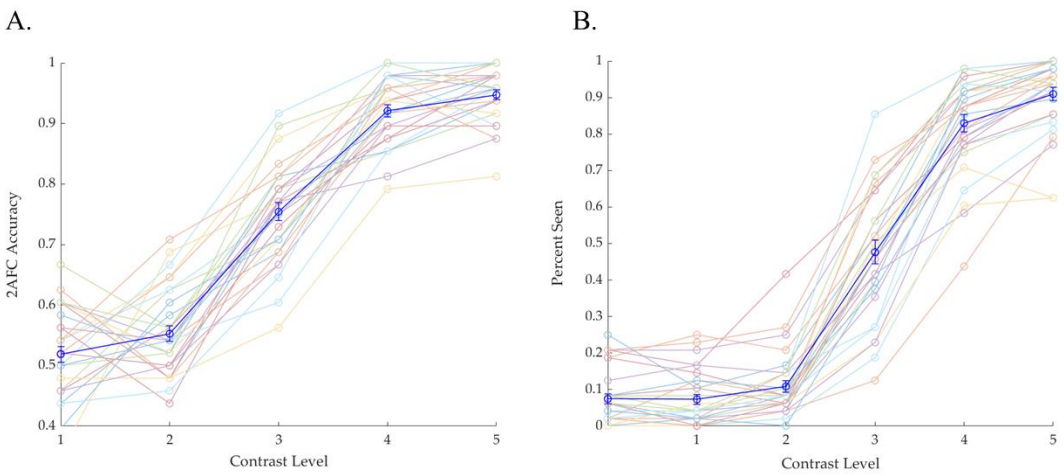

Figure 13: Behavioral functions from objective and subjective participant reports. Graphs show behavioral responses collected during the report condition. **A:** Participant accuracy on the 2AFC objective awareness task. Participants were instructed to indicate the orientation of the Gabor gratings on every trial. The solid dark blue line shows the mean performance across participants. Error bars represent standard error of the mean. Lightly colored lines show individual subject results. **B:** Participant visibility ratings on the PAS subjective awareness task. The solid dark blue line shows the mean proportion of trials reported as seen, as determined by a PAS rating of 2, 3, or 4. Error bars represent standard error of the mean. Lightly colored lines show individual subject results. Contrast level 0 indicates trials in which no stimulus was presented (blank control trials).

PAS rating responses were assessed by plotting the proportion of trials in which a Gabor grating rated as unseen (PAS rating of 1) versus seen (PAS rating of 2, 3, or 4). A similar pattern was observed for binary seen-versus-unseen PAS rating responses, with a slightly significant difference between levels 1 ($M=0.0726$, $SD=0.0753$) and 2 ($M=0.1091$, $SD=0.0869$, $t(29)=-3.1317$, $p=0.0039$), a significant difference between levels 2 and 3 ($M=0.4764$, $SD=0.1775$, $t(29)=-13.1300$, $p<0.0005$), a significant difference between levels 3 and 4 ($M=0.8299$, $SD=0.1310$, $t(29)=-14.2522$, $p<0.0005$), and a significant difference between levels 4

and 5 (M=0.9096, SD=0.1009, t(29)=-5.4496, p<0.0005). Again, an overall pattern of perceptual bifurcation was evident despite a quantitative difference between contrast level 1 and level 2, as well as contrast level 4 and level 5. Importantly, no significant difference was observed between blank trials (M=0.0745, SD=0.0753) and contrast level 1, suggesting that subjects reported seeing stimuli presented at the lowest contrast level the same as stimulus-absent trials. Bifurcation patterns observed for subjective ratings were confirmed by visualizing the proportion of each PAS rating per contrast level, which showed subjects consistently rated blank trials, contrast level 1 gratings, and contrast level 2 ratings a 1 or 2, and increased the use of 2, 3, and 4 for threshold, contrast level 4, and contrast level 5 gratings (Figure 14). Subjective awareness was also measured by plotting the average PAS rating given at each contrast level (Figure 15), complementing the pattern observed by plotting the proportion of PAS ratings at each contrast level. No significant difference in average PAS rating was observed between blank trials (M=1.0750, SD=0.0769) and contrast level 1 (M=1.0764, SD=0.0772, t(29)=-0.1483, p=0.8831), confirming binary seen versus unseen PAS responses to stimulus absent trials. Average PAS ratings increased in a non-linear manner as contrast level increased, with a slight significant difference observed between contrast level 1 and level 2 (M=1.1111, SD=0.0948, t(29)=-2.5603, p=0.0159) and significant differences observed between contrast level 2 and level 3 (M=1.5743, SD=0.2367, t(29)=-12.7215, p<0.0005), level 3 and level 4 (M=2.2826, SD=0.3721, t(29)=-16.1877, p<0.0005), and level 4 and level 5 (M=2.7583, SD=0.4857, t(29)=-12.1707, p<0.0005).

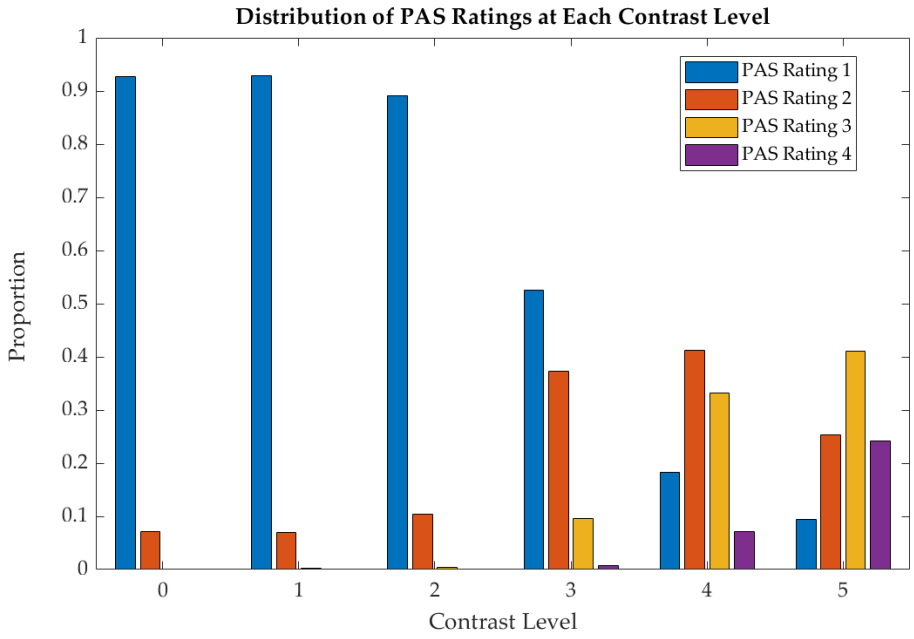

Figure 14: Distribution of PAS ratings at each contrast level.

PAS ratings were given by subjects during the report condition to measure subjective experience of the Gabor gratings. The proportion of trials for which a given rating (blue: PAS rating 1, orange: PAS rating 2, yellow: PAS rating 3, purple: PAS rating 4) was selected are plotted at each contrast level, including blank trials (contrast level 0).

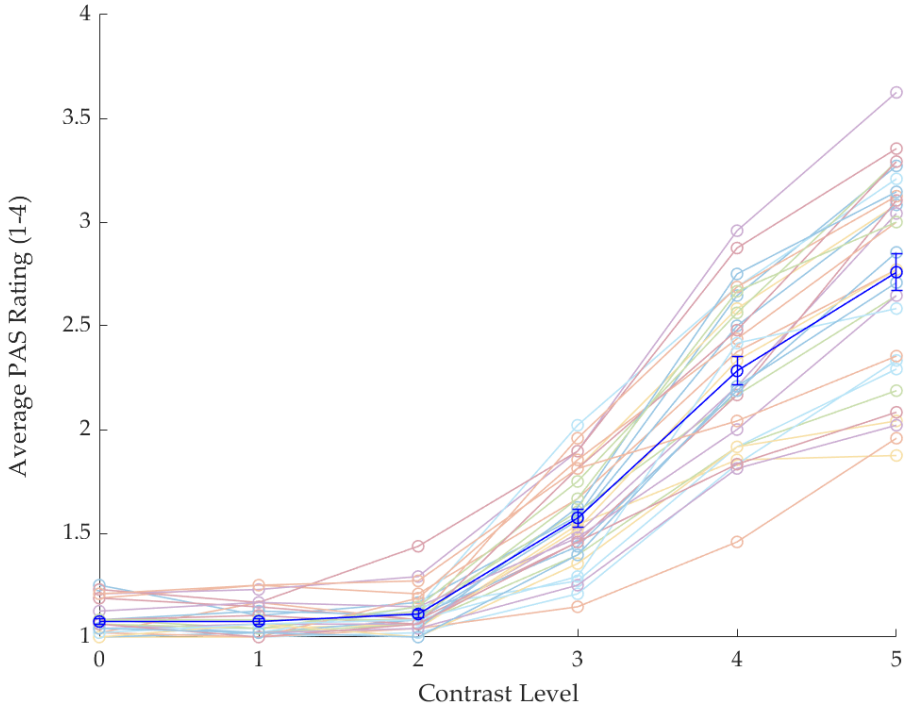

Figure 15: Average PAS Rating (1-4) at each contrast level.
PAS ratings were given by subjects during the report condition to measure subjective experience of the Gabor gratings, then averaged at each contrast level. The solid dark blue line shows the mean PAS rating across participants. Error bars represent standard error of the mean. Lightly colored lines show individual subject results.

Overall behavioral functions were tested against a linear and non-linear model using both linear and non-linear regressions (Figure 16). Mean square error (MSE) was used to assess whether behavioral data fit closest to a linear or non-linear sigmoidal model, with lower MSE values indicating better fit. Compared to linear regressions, MSE was lower for non-linear regressions performed on both the 2AFC accuracy (linear MSE=0.0786, non-linear MSE=0.0654) and binary seen-versus-unseen PAS ratings (linear MSE=0.1635, non-linear MSE=0.1129). Results from both the paired samples t-tests and the

regression analysis confirmed previous results demonstrating bifurcated patterns of perceptual report (Cohen et al., 2024; Del Cul et al., 2007; Sergent et al., 2021).

The average PAS rating behavioral function was also tested against linear and non-linear models to determine if different visualizations of subjective awareness yield qualitative differences (Figure 17). MSE was used to assess whether average PAS ratings fit closest to a linear (Figure 17A), exponential (Figure 17B), or sigmoidal model (Figure 17C), with lower MSE values again indicating better fit. Data best fit a sigmoidal model (MSE=0.2732) compared to exponential (MSE=0.2874) or linear (MSE=0.3646) models, though differences between the two non-linear models were minimal.

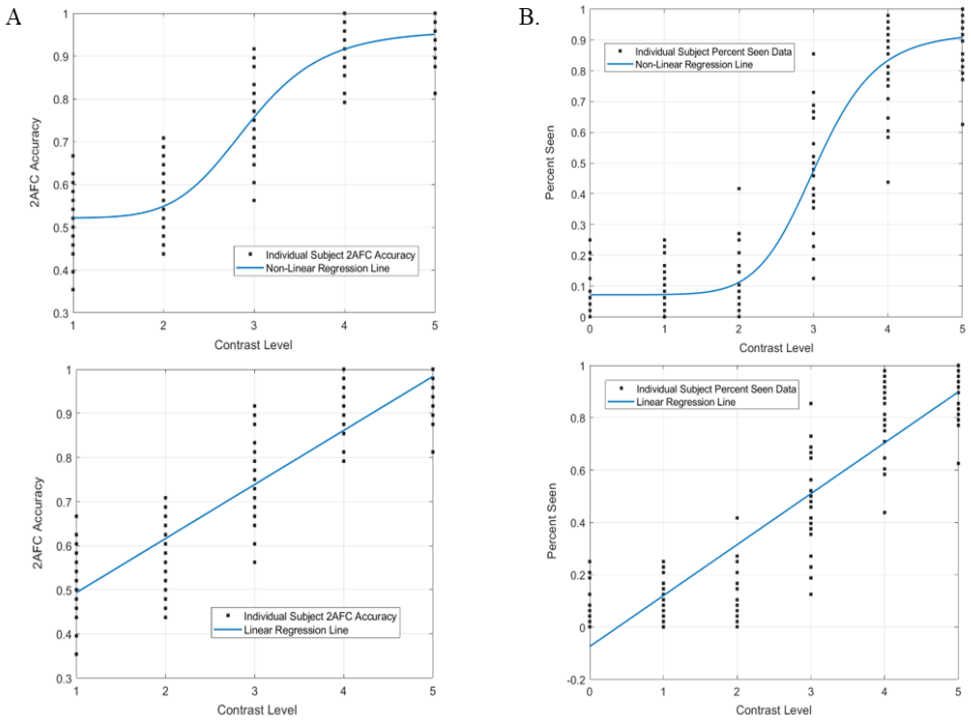

Figure 16: Linear and Non-Linear Regression Fitting of Behavioral Data.
Model fitting of objective (**A**) and subjective (**B**) awareness assessments collected during the report condition. Behavioral data were fit to linear (bottom) and non-linear (top) regression models. Data best fit a non-linear regression model, as demonstrated by lower MSE values compared to data fit with a linear model.

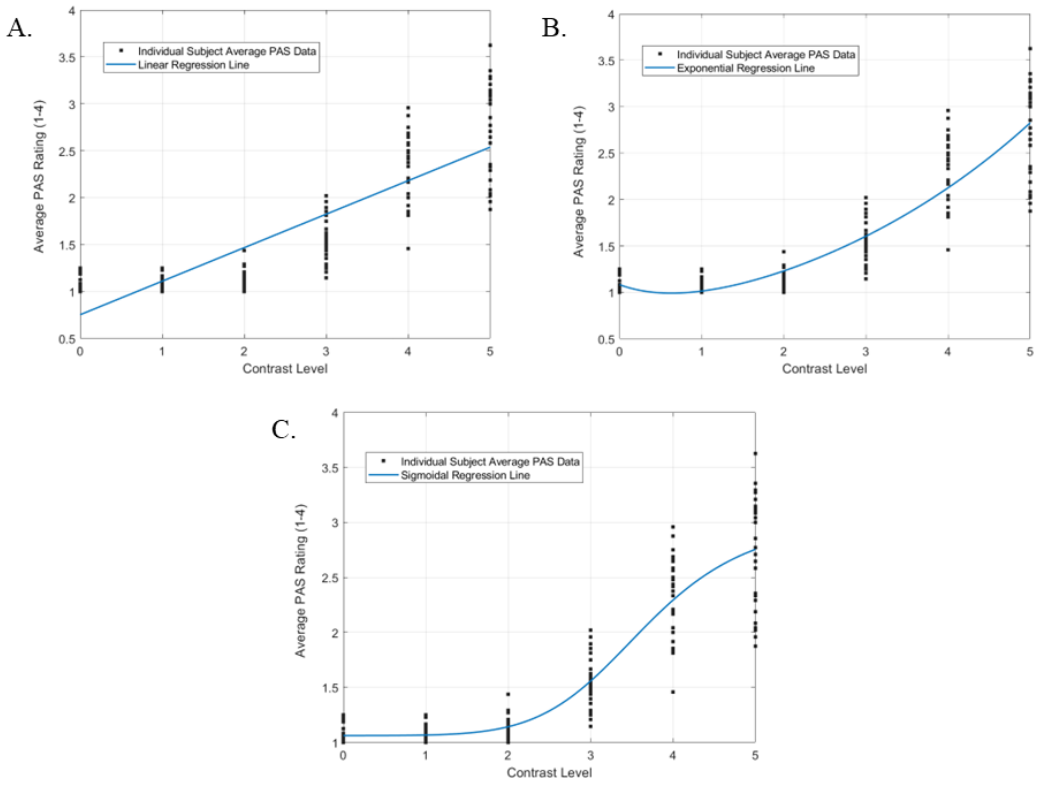

Figure 17: Linear and exponential fitting of average PAS ratings.
Average PAS ratings were fitted to a linear (**A**), exponential (**B**), and sigmoidal (**C**) model. Data best fit the sigmoidal regression model, as demonstrated by lower MSE values compared to data fit with a linear model.

3.2 Decoding Stimulus Presence at Each Contrast Level

Late metastable activity has been proposed as a valid signature of conscious processing (King and Dehaene, 2016; Dehaene 2014; Salti et al., 2015; Schurger et al., 2015). The majority of studies analyzing late patterns of metastability utilize report-based paradigms and contrastive methods of stimulus manipulation, leading to a possible over estimation of the neural

process related to conscious perception. We aimed to assess whether late metastable representations are observed for no-report paradigms utilizing parametric stimulus manipulation by implementing MVPA and temporal generalization of decoding. In the first analysis, multivariate pattern classifiers were trained to distinguish stimulus absent trials (i.e. blank trials) from stimulus present trials at each contrast level.

Same-time decoding plots (Figure 18) and temporal generalization matrices (Figure 19) were generated for each contrast level. Significant same-time decoding was absent below and at perceptual threshold, appearing only at contrast level 4 from 320-370ms and 550-600ms (peak AUC 56.90%) and contrast level 5 from 230ms to the end of the 1000ms epoch (peak AUC 67.15%). Classifiers were able to successfully distinguish between stimulus present and blank trials largely later in time, though above chance accuracy was observed earlier in time at contrast level 5 likely due to decreased variability between subjects and an improved signal-to-noise ratio.

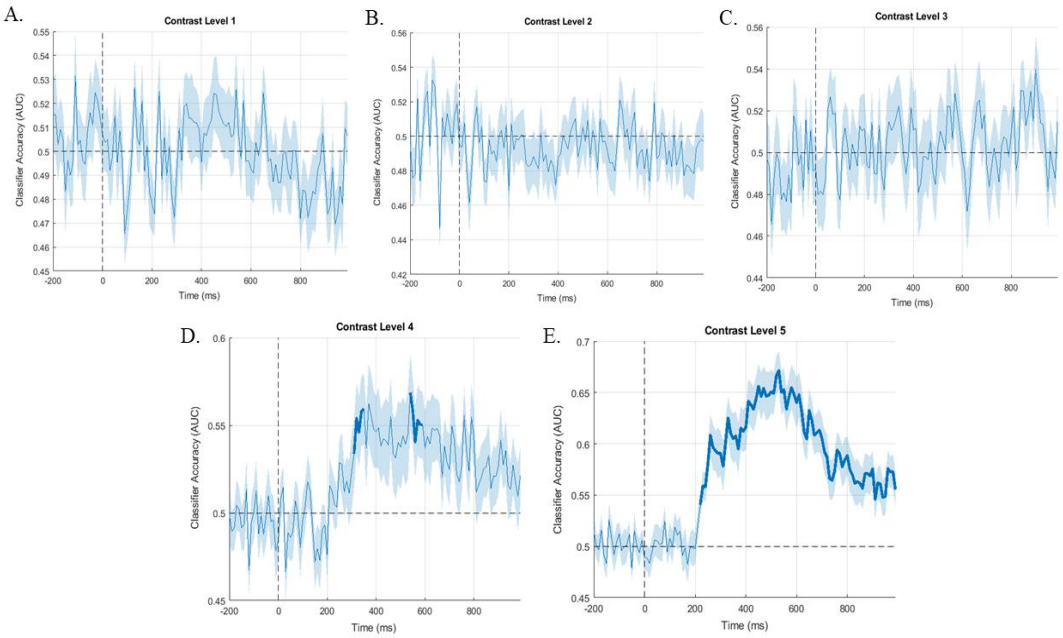

Figure 18: Same-time decoding of stimulus presence at each contrast level. Pattern classifiers were trained to distinguish between stimulus present and stimulus absent (i.e. blank) trials at each contrast level (**A**: level 1, **B**: level 2, **C**: level 3, **D**: level 4, **E**: level 5). Classification accuracy (AUC) is plotted over time (training and testing on the same time points), with thickened blue lines indicating significant above chance (50%) decoding after correction for multiple comparisons using cluster-based permutation tests ($p<0.05$). Shaded light blue regions represent standard error of the mean (SEM).

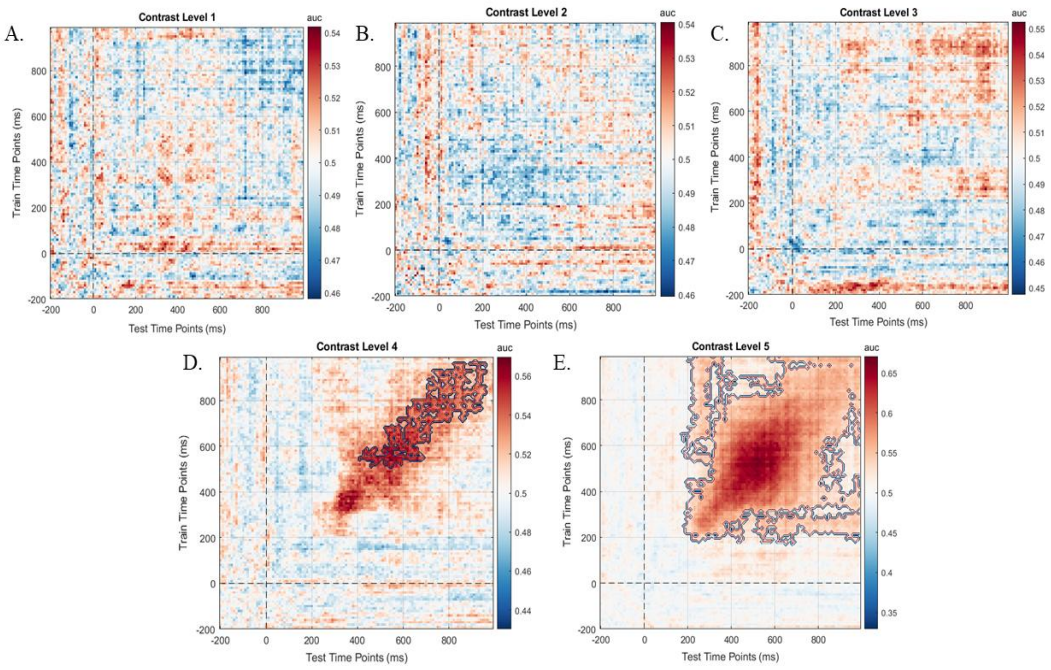

Figure 19: Temporal generalization of stimulus presence decoding at each contrast level.

Pattern classifiers were trained and tested at every time point to distinguish between stimulus present and stimulus absent (i.e. blank) trials at each contrast level (**A**: level 1, **B**: level 2, **C**: level 3, **D**: level 4, **E**: level 5). Classifier training time is plotted on the y-axis, and classifier testing time is plotted on the x-axis. Classification accuracy (AUC) is represented by color intensity of the heatmap, with significant above chance (50%) temporal generalization ($p<0.05$, as determined by cluster-based permutation testing) outlined.

Temporal generalization failed to demonstrate clear bifurcation dynamics, with no above chance decoding present at the lowest two contrast levels and threshold, and thick, square-shaped off-diagonal decoding that emerged at contrast level 4 and increased at contrast level 5. Significant off-diagonal decoding was observed at contrast level 3 from ~420-970ms (peak AUC 56.97%, $p<0.05$), though non-significant, slightly spread decoding was observed earlier in time starting at ~250ms. Generalization spread more robustly and earlier in time at contrast level 5, with significant decoding observed from ~180ms to the end of

the 1000ms epoch (peak AUC 67.08%, $p<0.05$). Two square-shaped patterns of temporal generalization were observed at the highest contrast level, representing two training-testing time windows of interest: an early time window from 200-400ms and a late time window from 400-600ms. These time windows were chosen for additional analysis based on previous studies identifying significant off-diagonal temporal generalization during periods both during and after the VAN (Hutchinson et al., 2024, 2025; Zhu et al., 2024).

To determine if temporal generalization exhibits bifurcation dynamics, classification accuracy of the early and late time windows (in both training and testing time) was plotted at each contrast level, then fit against linear, sigmoidal, and exponential (i.e. non-linear) regression models (Figure 20). Non-linear regression models were chosen based on bifurcation dynamics observed in the behavioral data, and by exponential amplitude increases as a function of visibility level observed in univariate ERP analysis. Classification accuracy in the early time window fit best with non-linear models (sigmoidal MSE=0.7819; exponential MSE=0.5135) compared to a linear model (MSE=1.6660). Results were similar for the late time window, data fitting non-linear models (sigmoidal MSE=0.4455; exponential MSE=0.3634) more closely than a linear model (MSE=3.7856). Stimulus presence decoding at each contrast level failed to exhibit precise bifurcation dynamics, instead showing a general pattern of non-linear classification accuracy increase. Non-linear patterns were observed during both early and late time windows, though lower MSE values overall were observed during the later time window. In other words, classifier accuracy remained low at contrast level 1, level 2, and level 3 (threshold), increased at contrast level 4, then substantially increased again at contrast level 5, showing a pattern that closely fits both sigmoidal and exponential models. The increase in accuracy observed between contrast levels 3 and 4 may be a result of perceptual threshold

increasing due to reduced attention to the Gabor gratings in no-report (see section 4.1). Overall, results from this initial decoding analysis indicate that brain activity encodes stimulus presence only when stimuli are perceived, suggesting a relationship between perceptual awareness and metastability.

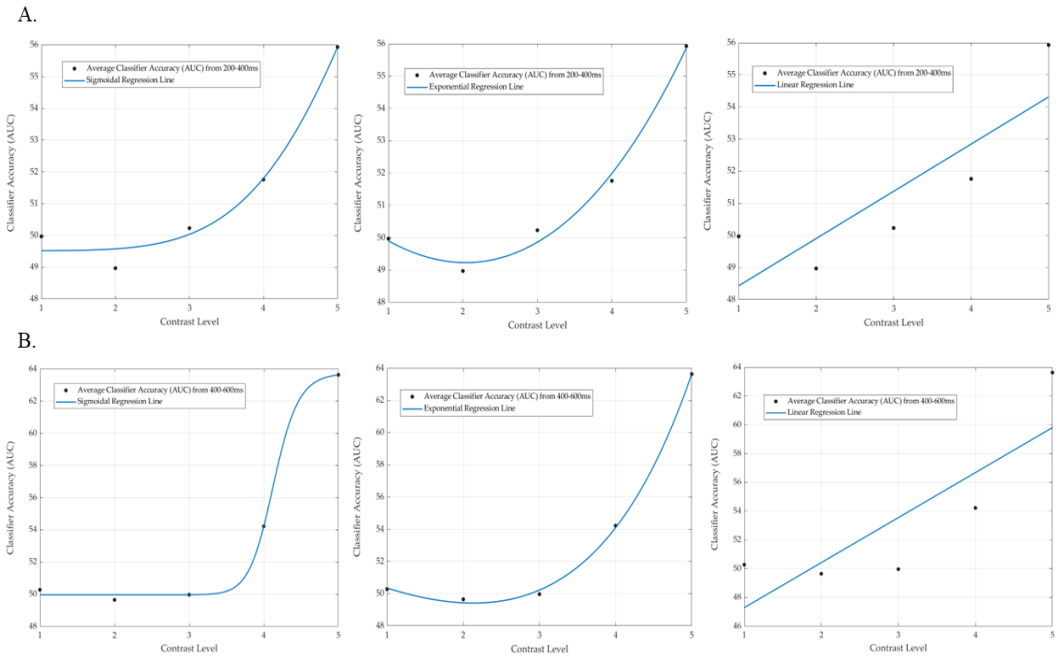

Figure 20: Linear and non-linear regression of average temporal generalization classifier accuracy for stimulus presence decoding at each contrast level. Classifiers were trained to distinguish between stimulus present and stimulus absent trials, then tested at all possible time points to generate a temporal generalization matrix. Classifier accuracy (AUC) was averaged for classifiers trained and tested from 200-400ms (**A**) and 400-600ms (**B**) at each contrast level, then fit to sigmoidal (right), exponential (middle), and linear (left) regression models.

3.3 Seen vs. Unseen Decoding

While decoding of stimulus presence appeared to suggest that temporal generalization increases with stimulus visibility, the extent of generalization when comparing seen and unseen stimuli remained unclear. To test this, classifiers were trained to distinguish Gabor gratings at contrast level 5 (most visible) from gratings at contrast level 1 (least visible), and tested at the same time points (same-time decoding, Figure 21A, top) and every other time point (temporal generalization, Figure 21A, bottom). Significant decoding along the same-time diagonal was observed from 240-830ms and 850-950ms (peak AUC 69.14%), suggesting distinct neural processing of seen versus unseen gratings throughout time. Similarly, temporal generalization showed robust spreading, with significant off-diagonal decoding observed from 160ms to the end of the 1000ms epoch (peak AUC 68.87%, $p<0.05$). Classifiers trained early in time (~200-400ms) were able to successfully decode between seen and unseen stimuli when tested later in time (~400-800ms), and vice versa, suggesting the neural processes encoding stimulus visibility remain stable over time. Results are consistent with results from stimulus presence decoding at contrast level 5, showing square-shaped generalization within an early (~200-400ms) and late time window (~400-600ms). Because stimuli at contrast level 1 are never perceived, and stimuli at contrast level 5 are always perceived, the observed significant spreading temporal generalization points to a relationship between metastability and conscious perception.

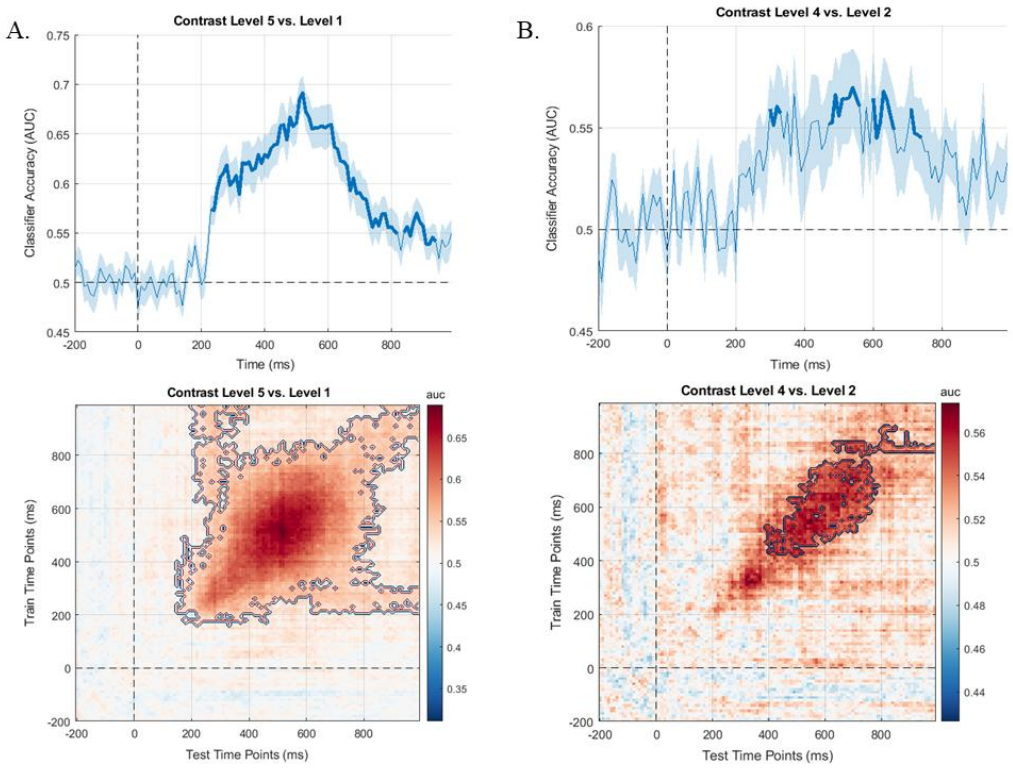

Figure 21. Decoding stimulus visibility contrast level 5 from level 1 (**A**) and level 4 from level 2 (**B**).

Classifiers were trained first to distinguish between Gabor gratings at contrast level 5 from gratings at contrast level 1 (**A**), visualized as both same-time decoding (top) and temporal generalization (bottom) plots. Thickened blue lines in the same-time decoding plots indicate significant above chance (50%) decoding after correction for multiple comparisons. Shaded light blue regions represent SEM. Temporal generalization classification accuracy (AUC) is represented by color intensity of the heatmap, with significant above chance (50%) temporal generalization ($p<0.05$). Classifiers were also trained to distinguish between Gabor gratings at contrast level 4 from gratings at level 2 (**B**), generating same-time decoding and temporal generalization plots (same as in panel **A**).

The previous analysis was modified to classify between contrast level 4 and level 2 Gabor gratings (Figure 21B) to determine if decoding remains robust for stimuli that differed less in physical contrast yet still resulted in always

unseen versus always seen reports. Significant same-time decoding (Figure 21B, top) was observed both early and later in time, albeit in more restricted time windows compared to level 5 versus level 1 classification (clusters identified from 310-340ms, 480-570ms, 610-670ms, 720-750ms, peak AUC 56.98%). Temporal generalization (Figure 21B, bottom) exhibited spreading along the diagonal, albeit less robust compared to decoding of contrast level 1 from level 5. Off-diagonal decoding was significant from ~400-800ms (peak AUC 57.36%, $p<0.05$). Above chance off-diagonal decoding was observed earlier in time (~200-400ms), though not significant, likely due to high variability between subjects and poor signal-to-noise ratios. The absence of above chance decoding before 200ms likely reflects a lack of differential neural activity related to physical stimulus differences, further validating the classification of contrast level 4 from level 2 as a seen versus unseen comparison.

The observed decrease in decoding accuracy and generalization spreading likely reflects the weaker differences between contrast level 4 and 2 compared to contrast level 5 and 1. Still, the source of these differences in generalization and classifier accuracy remain unclear, though, especially given the lack of perceptual differences between contrast levels 1 and 2 (both almost never seen) and levels 4 and 5 (both almost always seen). These results suggest that temporal generalization does not display clear bifurcation dynamics, as neural differences are present where perceptual differences are absent.

3.4 Decoding Adjacent Contrast Levels

To further parse the relationship between temporal generalization, perceptual awareness, and bifurcation dynamics, classifiers were trained to distinguish Gabor gratings between adjacent contrast levels (i.e. distinguish a

Gabor grating at contrast level 3 from a Gabor grating at contrast level 4). As with the previous analysis, both same-time decoding plots (Figure 22) and temporal generalization matrices (Figure 23) were generated. Bifurcation dynamics of perception may be reflected by adjacent contrast level decoding if no above chance decoding is observed between minimally and maximally visible contrast levels (i.e. between levels 1 and 2 and levels 4 and 5) and if significant above chance decoding is observed between threshold and surrounding contrast levels (i.e. between levels 2 and 3 and levels 3 and 4). In other words, above chance decoding accuracy between contrast levels in which changes are related only to physical stimulus differences likely reflects unconscious sensory processing (particularly if decoding is early in time and restricted only to the same-time diagonal). Conversely, above chance decoding between contrast levels in which there are both sensory and perceptual changes likely reflects neural activity associated with conscious processing. Ultimately, decoding of adjacent contrast levels likely reflects the accessibility of content to conscious perception, which maintains important implications for the debate between sensory and cognitive theories of consciousness (see section 1.5).

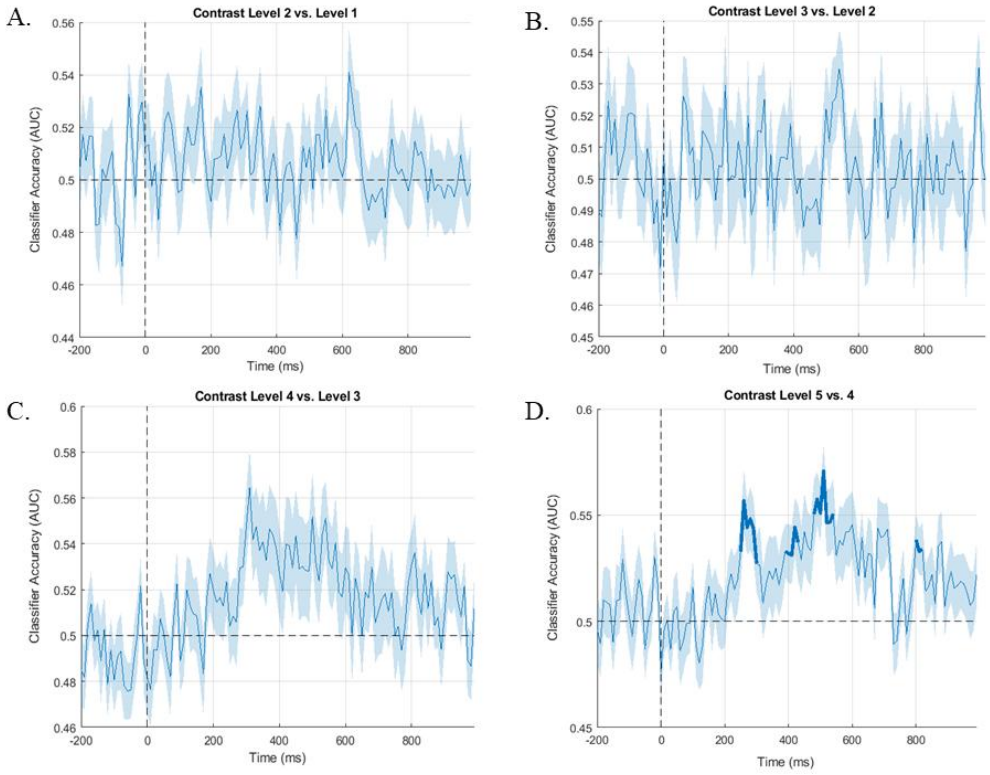

Figure 22. Same-time decoding of adjacent contrast levels.
Classifiers were trained to distinguish between Gabor gratings at adjacent contrast levels, then tested at the same time points to generate same-time decoding plots (**A**: level 1 vs. level 2, **B**: level 2 vs. level 3, **C**: level 3 vs. level 4, **D**: level 4 vs. level 5). Thickened blue lines indicate significant above chance (50%) decoding after correction for multiple comparisons. Shaded light blue regions represent SEM. Significant above chance decoding was only observed for classifiers trained to distinguish between level 4 and level 5.

Significant same-time diagonal decoding was only observed for classifiers trained to distinguish gratings at contrast level 5 from gratings at contrast level 4 (Figure 22D), with above chance accuracy both early and later in time (clusters identified from 260-310ms, 400-440ms, 490-550ms, and 810-830ms, peak AUC 57.08%). Non-significant above chance decoding was also observed for classifiers trained to distinguish between contrast levels 4 and 3, as well as contrast levels 3

and 2. Classifier accuracy peaked for the former classification ~300-550ms (peak AUC 56.45%), closely aligning with the significant decoding time windows observed when classifying contrast level 5 versus level 4. While overall classification accuracy decreased when classifying level 3 versus level 2, a peak emerged from ~500-600ms, again aligning with late above chance decoding observed for contrast level 5 versus level 4 classification. The non-significant above chance decoding peaks observed when classifying between contrast levels adjacent to perceptual threshold may represent differential neural activity related to perception, yet lack significance due to poor signal-to-noise ratios and/or excessive variability between subjects. Conversely, the observed peaks may be unrelated to perceptual differences, and be solely reflective of noise in the data.

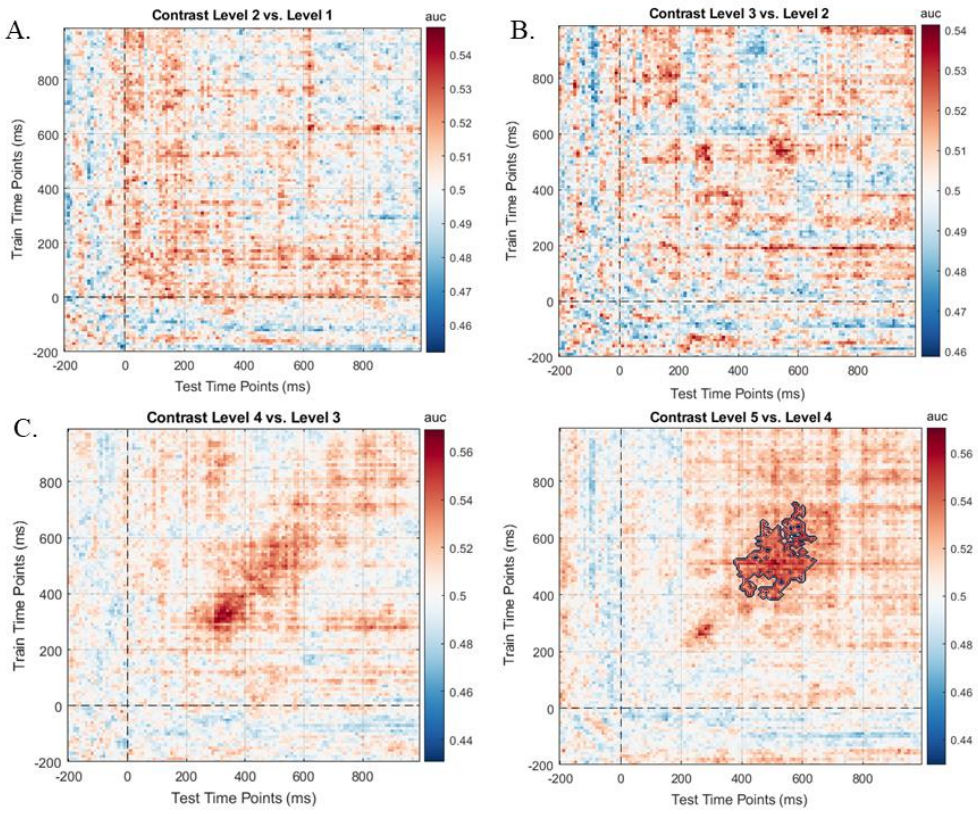

Figure 23. Temporal generalization of adjacent contrast level decoding. Classifiers were trained to distinguish between Gabor gratings at adjacent contrast levels, then tested at all time points to generate temporal generalization matrices (**A**: level 1 vs. level 2, **B**: level 2 vs. level 3, **C**: level 3 vs. level 4, **D**: level 4 vs. level 5). Classifier training time is plotted on the y-axis, and classifier testing time is plotted on the x-axis. Classification accuracy (AUC) is represented by color intensity of the heatmap, with significant above chance (50%) temporal generalization ($p<0.05$, as determined by cluster-based permutation testing) outlined. As with same-time decoding (Figure 19), significant above chance decoding was only observed for classifiers trained to distinguish level 4 from level 5, with limited off-diagonal spreading.

Temporal generalization deviated from our expected results, showing no significant spreading between contrast levels 2 and 3 or contrast levels 3 and 4 (Figure 23A, B, and C). Significant above chance decoding was only observed for level 4 vs. level 5 decoding, with limited off-diagonal spreading from ~390-650ms

(peak AUC 57.05%). Again, temporal generalization failed to display bifurcation dynamics, demonstrated by the lack of significant off-diagonal decoding between contrast levels around threshold (level 2 vs. level 3 and level 3 vs. level 4). What the limited spreading observed for level 4 vs. level 5 decoding represents remains unclear. 2AFC and binary split (i.e. counting trials with a rating of 1 as unseen and trials with a 2 or higher as seen) PAS data indicate no perceptual differences between Gabor gratings at contrast level 4 and contrast level 5, yet the relatively late timing of the significant off-diagonal decoding suggests neural differences are likely not solely related to physical stimulus differences. This pattern may instead be explained by average PAS rating data, which show both a quantitative and qualitative difference in subjective experience between contrast level 4 and 5. While no significant decoding was observed when classifying level 3 vs. level 4, classifiers exhibited relatively high accuracy and limited spreading of generalization earlier in time from ~250-400ms (peak AUC 56.93%). Taken together, these results may demonstrate slight neural differences earlier in time (i.e. within the time window of the VAN) related to perception between threshold and contrast level 4, and more robust neural differences later in time (i.e. after the VAN time window).

3.5 Decoding Stimulus Orientation at Each Contrast Level

After decoding the accessibility of conscious content, we sought to test whether classifiers could successfully decode properties of the conscious content, in seen and unseen conditions separately. To do so, pattern classifiers were trained to distinguish right oriented Gabor gratings (45˚) from left oriented

Gabor gratings (135°) at each contrast level, then tested at the same time points (same-time decoding) and every other time point (temporal generalization of decoding). In line with previous analyses, we predicted that, if temporal generalization displays bifurcation dynamics, no significant above chance decoding would be observed at contrast level 1 and level 2, limited significant decoding would be present at threshold, and more robust temporal generalization of decoding would be present at levels 4 and 5.

Surprisingly, no above chance decoding was observed across all contrast levels. Decoding largely remained at or below chance levels, even as stimulus visibility increased. Pilot data from the report condition were also analyzed, yet decoding remained at chance even at the highest contrast level. These results are inconsistent with previous studies, which show successful above chance decoding of Gabor grating orientation from scalp EEG data (Bae & Luck, 2018; Nguyen et al., 2020; Ramkumar et al., 2013).

Chapter 4: Discussion

Over the last 35 years, one question has plagued consciousness researchers: what is the best way to isolate neural activity related to conscious perception? The current study sought to answer this question by implementing a no-report parametric stimulus degradation paradigm to link bifurcated perceptual reports to scalp EEG data. Data were analyzed using both univariate (ERPs) and multivariate (decoding and temporal generalization of decoding, the focus of this thesis) methods, allowing for a more robust understanding of conscious processing at and around perceptual threshold. By parametrically manipulating stimulus visibility in the absence of report, preconscious and post-perceptual processes were more effectively separated from potential NCCs.

Bifurcation dynamics have been proposed as a potential signature of conscious processing (Cohen et al., 2024; Del Cul et al., 2007; Sergent et al., 2021), but evidence is limited, particularly for no-report visual paradigms. To our knowledge, only Sergent et al. (2021) and Cohen et al. (2024) have examined neural bifurcation dynamics with MPVA (same-time decoding and temporal generalization). Both studies failed to identify bifurcation dynamics of decoding under no-report conditions, challenging previous studies highlighting late temporal generalization as a potential NCC (Hutchinson et al., 2025; Dehaene & King, 2016; King et al., 2016; King & Dehaene, 2014). The answer to whether late metastability is a viable candidate NCC may also help substantiate or refute claims made by both sensory and cognitive theories of consciousness.

4.1 Temporal Generalization as a Potential NCC

Our primary aim was to determine if late temporal generalization, or metastability, displays bifurcation dynamics in line with objective task performance and binary subjective awareness behavioral measures, and is thus a potential NCC. Behavioral reports collected during the report condition demonstrated clear bifurcation dynamics, where stimuli below perceptual threshold were almost never seen, and stimuli above perceptual threshold were almost always seen. Bifurcation dynamics of perception were present for both objective and subjective reports, suggesting a limited effect of unconscious processing on the identification of physical stimulus features (e.g., Gabor grating orientation) below perceptual threshold. Separating PAS responses into binary seen versus unseen categories confirmed a qualitative change in perception unrelated to increases in physical stimulus strength for stimuli above threshold, in line with results from previous studies (Cohen et al., 2024; Del Cul et al., 2007; Sergent et al., 2021). Average PAS ratings at each contrast level also demonstrate a qualitative change in perception, but diverge from binary PAS measures in that consistent changes in perception are present once stimuli are perceived (i.e. contrast levels 3, 4, and 5). Both measurements of subjective experience suggest stimuli below threshold are almost never perceived, and that stimuli above threshold are always seen but are experienced at qualitatively increasing levels of clarity as physical contrast continues to increase. Behavioral data generally point to a clear boundary between unconscious and conscious processing of stimuli, though, allowing potential NCCs to be compared with each of the observed psychometric functions to confirm whether the signal indexes conscious processing, or some other pre-conscious or post-perceptual process.

Clear bifurcation dynamics were absent when decoding stimulus present from stimulus absent trials, though classification accuracy increased in a largely non-linear manner, similar to average PAS rating behavioral measurements. Still, interesting temporal generalization patterns emerged that may reflect the transition from unconscious to conscious processing. In line with behavioral reports, above chance decoding was absent below perceptual threshold, suggesting no distinct neural differences between noise-only blank trials and low contrast gratings. The lack of same-time diagonal decoding below threshold highlights a possible lack of distinguishable unconscious sensory processing due to extremely minimal stimulus differences, particularly when compared to results from previous studies employing much more dramatic manipulations of stimulus visibility such as backwards masking (Cohen et al., 2024; Del Cul et al., 2007). Overall decoding accuracy slightly increased for stimuli at threshold-levels of contrast, though no significant patterns of generalization were observed. Evidence of non-significant generalization emerged later in time, showing a minimally square-shaped pattern possibly related to conscious perception. High variability between subjects, along with poor signal-to-noise likely limited classification accuracy at threshold.

At contrast level 4, significant temporal generalization was observed only later in time, with limited off-diagonal spreading. Non-significant generalization showed more robust spreading throughout time, including during early time windows associated with processing of perceptual content (i.e. the VAN). Decoding at contrast level 5 demonstrated high classification accuracy and spread, with square-shaped patterns of generalization both early and late in time. The square-shaped temporal generalization patterns observed at contrast levels 4 and 5 appear similar to results from report condition decoding (Cohen et al., 2024; Hutchinson et al., 2025; Sergent et al., 2021; Zhu et al., 2024), which

typically shows significant spreading of off-diagonal decoding both early and later in time. While these results are in line with proposed patterns of late metastability related to conscious processing, they have more recently been challenged as possibly reflecting post-perceptual processes rather than conscious processing alone (Cohen et al., 2024; Hutchinson et al., 2025; Zhu et al., 2024). The observed differences between contrast level 5, contrast level 4, and threshold, then, may be the result of post-perceptual processes that lingered even in the absence of perceptual report. Alternatively, the significant square-shaped temporal generalization of decoding observed at contrast level 5 may genuinely reflect conscious processing of seen Gabor gratings in no-report, with minimal task-related processing confounds. Instead, the differences in classification accuracy and robustness of spreading observed between threshold, contrast level 4, and contrast level 5 may be the result of poor signal-to-noise ratios and/or high variability between subjects as a consequence of the simplicity of the stimuli and minimal physical stimulus differences created by embedding the stimuli in continuous visual static. The lack of clear bifurcation dynamics may also be the result of reduced attention to stimuli in no-report, reducing perceived visibility of the gratings and thus shifting the true perceptual threshold to a higher level of visual contrast. In other words, rendering stimuli task-irrelevant may have reduced their overall visibility, resulting in stimuli at threshold being perceived less than 50% of the time. This would both explain the minimal temporal generalization observed at threshold, as well as the apparently shifted sigmoidal function yielded when comparing classifier performance at each contrast level, particularly later in time.

Stimulus presence decoding results may be better explained by comparing classifier accuracy at each contrast level to the average PAS rating data, as both show non-sigmoidal, non-linear patterns of increase. Decoding during the early

time window matches the average PAS data particularly well, with differences in subjective experience and classifier accuracy present between contrast level 4 and 5. While the functions do not perfectly match, they point to a possible relationship between stimulus strength and the quality of perception once stimuli are above perceptual threshold. While objective task performance and binary subjective reports reflect whether a stimulus was perceived or not, perhaps average PAS rating data points to more fine-grained conscious processing through which subjective experience of stimuli become more phenomenologically clear and detailed as physical stimulus strength increases.

Despite not displaying clear bifurcation dynamics, stimulus presence decoding above perceptual threshold ultimately points to markers of conscious processing later in time, particularly after the typical time window of earlier ERP signals like the VAN (Demski et al., 2021; Förster et al., 2020; Koch et al., 2016). A handful of potential ERP candidate NCCs have been identified beyond the typical VAN time window (~100-300ms), including the fronto-central N2 (Cogitate et al., 2023; Cohen et al., 2024 ; Cohen et al., 2020; Dellert et al., 2022; Gaillard et al., 2009; Schlossmacher et al., 2020; Sergent et al., 2021) and late posterior positivity (Cohen et al., 2020; Cohen et al., 2024; Dellert et al., 2021; Pitts et al., 2014; Sergent et al., 2021; Shafto & Pitts, 2015). ERP analysis of the current study revealed a late posterior positive signal exhibiting non-linear, albeit not sigmoidal, increases in amplitude as a function of stimulus visibility from 450-650ms, a similar time window to above threshold stimulus presence temporal generalization of decoding. While the late posterior positivity or temporal generalization of stimulus presence decoding fail to clearly bifurcate, either may still represent a later stage of conscious processing. These results are still unclear, though, as no significant above chance decoding was observed during earlier

time windows, even along the same-time diagonal, as was found by Cohen et al. (2024).

ERP analysis of the current study also revealed an occipital negativity from 250-350ms, identified as the VAN. In depth analysis of the VAN revealed non-linear, albeit not sigmoidal (i.e. not matching objective performance and binary PAS subjective report data), increases in amplitude. While the VAN does not display clear bifurcation dynamics, clear amplitude increases are present, between threshold and contrast levels 4 and 5. These changes in amplitude align closely with the off-diagonal, slightly square-shaped decoding observed when classifying stimulus presence at contrast levels 4 and 5, suggesting early brain signals may be necessary for consciousness. No differences in VAN amplitude were found between contrast level 1, contrast level 2, and threshold, further corroborating the current stimulus presence decoding results.

Additional decoding analyses sought to further relate neural activity to behavioral responses and validate results from initial stimulus presence classification. Classifiers were successfully able to distinguish Gabor gratings at contrast level 5 from gratings at contrast level 1, resulting in robust, square-shaped patterns of temporal generalization within and beyond the time window of the VAN. Results were nearly identical to contrast level 5 stimulus presence decoding, confirming no perceptual or neural differences between blank trials and gratings at contrast level 1. These results closely mirror traditional contrastive seen versus unseen comparisons, which have demonstrated similar square-shaped spreading both early and later in time, but the results here suggest that such patterns may persist even in the absence of report (Hutchinson et al., 2024, 2025; Sergent et al., 2021; Zhu et al., 2024). A similar pattern of temporal generalization was observed for level 4 versus level 2 decoding, though less robust, as was observed for contrast level 4 stimulus presence classification.

While generalization was more restricted to the same-time diagonal, spreading still appeared throughout time, including during and after the VAN time window. These results likely reflect the more subtle physical stimulus differences between contrast level 4 and level 2 compared to contrast level 5 and level 1 rather than perceptual differences between adjacent contrast levels above and below threshold (i.e. levels 1 and 2 or levels 4 and 5). Both stimulus presence decoding and seen versus unseen decoding point to temporal generalization early and late in time as a marker of conscious perception.

Results from adjacent contrast level decoding support findings from stimulus presence and seen vs. unseen decoding. In line with average PAS rating subjective measurements, classifiers were able to successfully distinguish between gratings at contrast level 5 from gratings at contrast level 4 during early and late time windows, similar to those observed in the previous two analyses. Non-significant off-diagonal decoding was observed when classifying contrast level 4 from level 3, with classification accuracy peaking during the VAN time window. Additionally, no clear above chance decoding was observed when classifying level 3 from level 2, despite subjects reporting clear perceptual differences (albeit in the report condition with increased attention). Level 2 versus level 1 classification also aligned with average PAS ratings, showing no above chance decoding between perceptually identical stimuli. While these results fail to match objective task performance and binary subjective awareness measures, they do not fully contradict previous results or the possibility that metastability indexes conscious perception, especially given more fine-grained behavioral measurements of subjective awareness, such as the more continuous PAS.

4.2 Theoretical Implications

The precise timing of conscious processing, and the location of underlying neural generators, remains highly debated. Sensory and cognitive theories of consciousness make largely distinct predictions regarding the potential time windows and brain regions implicated in conscious processing, which may be reflected by the duration and pattern of temporal generalization observed in a given experiment. Taken together, the current results support neither sensory nor cognitive theories of consciousness. Stimulus presence and seen versus unseen decoding results, for instance, align closest with GNWT, which predicts late sustained temporal generalization in response to consciously perceived stimuli (King & Dehaene, 2014; Dehaene & King, 2016). These metastable representations are a result of sensory information being broadcast to the global workspace, where information can then be accessed by cognitive networks. The late, spread temporal generalization observed when decoding stimulus visibility, then, may reflect this process, though without the accessing of sensory information from cognitive networks due to an absence of perceptual report. The presence of early, spread generalization for all decoding analyses, on the other hand, also aligns with sensory theories of consciousness, which predict temporal generalization in posterior sensory cortices lasting only for the duration of conscious experience (Melloni et al., 2021, 2023). The limited square-shaped generalization pattern observed during the time window of the VAN may reflect a process similar to this account, though the duration of generalization persisted longer than the duration of the stimulus (i.e. longer than 100ms).

A larger question remains regarding the effect of task relevance on the timing and extent of temporal generalization. Several studies have recently sought to refute GNWT citing an absence of the P3b in no-report paradigms.

These studies argue that late signals reflect post-perceptual processes outside of conscious processing, suggesting candidate NCCs exist either only early in time, or are unrelated to global broadcasting. Few studies have successfully demonstrated late temporal generalization of decoding in no-report paradigms, adding to speculation that late metastability may be related to task-relevance rather than conscious processing (as with the P3b). New evidence from a handful of studies challenge the assertion that the P3b and metastability are linked, though, demonstrating fairly robust temporal generalization in no-report paradigms (Hutchinson et al., 2025; Sergent et al., 2021; Zhu et al., 2024). Temporal generalization in these studies, while present, occurs earlier in time than predicted by GNWT, spanning ~200-600ms and exhibiting thick diagonal rather than square-shaped patterns. Based on these results, only brief metastability may be required for conscious perception, with later periods of metastability required only for cognitive and other post-perceptual processes. More work is needed to confirm this possibility, however, especially given the limited number of studies demonstrating significant temporal generalization in no-report paradigms. Notably, Cohen et al. (2024) failed to observe temporal generalization in the no-report condition, finding only diagonal decoding later in time for stimuli above perceptual threshold. Metastability may still be a unique marker of conscious processing, as these findings are likely due to a suppression of sensory information maintenance by the mask stimuli, which were always perceived.

If brief metastability during intermediate time periods is a true marker of conscious processing, then claims made by both sensory and cognitive theories of consciousness hold weight. In line with sensory theories, early temporal generalization during the time window of the VAN may index the stability of conscious content. Conversely, later temporal generalization may index the

accessibility of the stabilized conscious content, resulting in further metastability during time windows after the VAN and extending beyond ~600ms only when content is actively used by post-perceptual networks. The current results support this two-stage unification of cognitive and sensory theories of consciousness, showing early, more restrained square-shaped generalization early in time, then more robust square-shaped spreading later in time, all in the absence of behavioral report. The non-linear increase in classifier accuracy observed for stimulus presence decoding also supports this view, possibly demonstrating the process of sensory evidence accumulation leading to perceptual content being encoded in an accessible format, then that content being accessed when above perceptual threshold.

4.3 Limitations and Future Directions

The current study suffers from multiple limitations, paving the way for future research into temporal generalization as a potential NCC. The primary limitation of this study centers around the use of weak Gabor grating stimuli embedded in continuous dynamic visual noise. While the use of simple stimuli allowed for a more streamlined experimental design and isolation of potential NCCs not specifically related to more complex visual processing (i.e. processing of faces, or processing of a stimulus quickly followed by a pattern mask), EEG responses maintained poor signal-to-noise ratios that hindered clear analysis of contrast levels directly below and above threshold. The simple Gabor stimuli also elicited reduced brain responses overall, with ERP amplitudes reaching only ±2μV at the highest contrast level in no-report. Together, low overall amplitude and poor signal-to-noise ratios prevented clear analysis of processing around

perceptual threshold, as well as proper assessment of possible bifurcation dynamics of temporal generalization.

Previous no-report studies focus on manipulating the awareness of more salient stimuli to isolate potential correlates of consciousness, highlighting additional limitations of parametric stimulus degradation. Inattentional blindness (Hutchinson et al., 2024, 2025; Pitts et al., 2012, 2014; Shafto & Pitts, 2015), masking (Cohen et al., 2020, 2024), and sine-wave speech studies (Zhu et al., 2024), for instance, always use highly salient stimuli, though with masking those stimuli are followed by an even more salient mask and with sine-wave stimuli speech content is degraded. For these types of awareness manipulations, it may be easier to assume that perception of stimuli above perceptual threshold is approximately the same in no-report compared to report, as stimuli are highly salient in both conditions. The current study, then, perhaps demonstrates a methodological limit for applying no-report conditions to different manipulations of awareness. That is, studies employing stimulus degradation may always suffer from attentional effects on the placement of perceptual threshold, cause perceptual threshold to shift to higher levels of physical stimulus strength in no-report compared to report. Conversely, inattentional blindness, masking, and possibly dichoptic color fusion (Schurger et al., 2025) allow for seen-versus-unseen contrasts that are perceptually similar in report and no-report conditions.

Another major limitation is the lack of a method for quantitatively assessing the extent of off-diagonal generalization. Currently, there are no agreed-upon analysis methods for assessing how much time generalization off the same-time diagonal constitutes true metastability rather than noise. While the current results consider ~200ms of generalization (in training and testing time) to

be metastability, more quantitative methods are necessary to fully confirm this claim.

Future studies should address these limitations by modifying the current experimental design to use more complex stimuli (e.g., line-drawings of objects) embedded in static visual noise, increasing the overall amplitude of brain responses and improving signal-to-noise ratios. Alternatively, a larger span of stimulus strength (e.g., including higher contrast levels) might be needed when no-report conditions are combined with stimulus degradation paradigms to compensate for the likely rightward-shifting of the perceptual threshold when task-based attention is withdrawn from the stimuli. Increasing the span of stimulus strength may also be paired with a more continuous subjective awareness scale rather than the typical 4-point PAS scale to achieve more fine-grained behavioral measures. Additional work should also be done to quantitatively assess temporal generalization, focusing on what specific stimulus manipulations, types, and durations elicit true metastable representations in no-report. Such studies may also incorporate more spatially sensitive imaging methods such as fMRI to better determine what brain regions are implicated in metastable processes.

Bibliography

B., M., & Boring, E. G. (1950). A History of Experimental Psychology. The American Journal of Psychology, 63(4), 623–631. https://doi.org/10.2307/1418885

Bae, G.-Y., & Luck, S. J. (2018). Dissociable Decoding of Spatial Attention and Working Memory from EEG Oscillations and Sustained Potentials. The Journal of Neuroscience, 38(2), 409–422. https://doi.org/10.1523/JNEUROSCI.2860-17.2017

Block, N. (2002). The Harder Problem of Consciousness. The Journal of Philosophy, 99(8), 391–425. https://doi.org/10.2307/3655621

Block, N. (2019). What Is Wrong with the No-Report Paradigm and How to Fix It. Trends in Cognitive Sciences, 23(12), 1003–1013. https://doi.org/10.1016/j.tics.2019.10.001

Boly, M., Massimini, M., Tsuchiya, N., Postle, B. R., Koch, C., & Tononi, G. (2017). Are the Neural Correlates of Consciousness in the Front or in the Back of the Cerebral Cortex? Clinical and Neuroimaging Evidence. The Journal of Neuroscience, 37(40), 9603–9613. https://doi.org/10.1523/jneurosci.3218-16.2017

Brascamp, J., Blake, R., & Knapen, T. (2015). Negligible fronto-parietal BOLD activity accompanying unreportable switches in bistable perception. Nature Neuroscience, 18(11), 1672–1678. https://doi.org/10.1038/nn.4130

Brown, R., Lau, H., & LeDoux, J. E. (2019). Understanding the Higher-Order Approach to Consciousness. Trends in Cognitive Sciences, 23(9), 754–768. https://doi.org/10.1016/j.tics.2019.06.009

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

Charles, L., King, J.-R., & Dehaene, S. (2014). Decoding the Dynamics of Action, Intention, and Error Detection for Conscious and Subliminal Stimuli. The Journal of Neuroscience, 34(4), 1158. https://doi.org/10.1523/JNEUROSCI.2465-13.2014

Cogitate, C., Ferrante, O., Gorska-Klimowska, U., Henin, S., Hirschhorn, R., Khalaf, A., Lepauvre, A., Liu, L., Richter, D., Vidal, Y., Bonacchi, N., Brown, T., Sripad, P., Armendariz, M., Bendtz, K., Ghafari, T., Hetenyi, D., Jeschke, J., Kozma, C., … Melloni, L. (2023). An adversarial collaboration to critically evaluate theories of consciousness (p. 2023.06.23.546249). bioRxiv. https://doi.org/10.1101/2023.06.23.546249

Cohen, M. A., Dembski, C., Ortego, K., Steinhibler, C., & Pitts, M. (2024). Neural signatures of visual awareness independent of postperceptual processing. Cerebral Cortex, 34(11), bhae415. https://doi.org/10.1093/cercor/bhae415

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

Cul, A. D., 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

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

Dehaene, S., & King, J.-R. (2016). Decoding the Dynamics of Conscious Perception: The Temporal Generalization Method. In Micro-, Meso- and Macro-Dynamics of the Brain. Springer; eBook Open Access (OA) Collection (EBSCOhost). https://search.ebscohost.com/login.aspx?direct=true&db=e001mww&AN=3734214&site=ehost-live

Dehaene, Stanislas. (2014). Consciousness and the brain: Deciphering how the brain codes our thoughts. Viking.

Dellert, T., Krebs, S., Bruchmann, M., Schindler, S., Peters, A., & Straube, T. (2022). Neural correlates of consciousness in an attentional blink paradigm with uncertain target relevance. NeuroImage (Orlando, Fla.), 264, 119679–119679. https://doi.org/10.1016/j.neuroimage.2022.119679

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

Delorme, A., & Makeig, S. (2004). EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. Journal of Neuroscience Methods, 134(1), 9–21. https://doi.org/10.1016/j.jneumeth.2003.10.009

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

Duman, I., Ehmann, I. S., Gonsalves, A. R., Gültekin, Z., Van den Berckt, J., & van Leeuwen, C. (2022). The No-Report Paradigm: A Revolution in Consciousness Research? Frontiers in Human Neuroscience, 16. https://doi.org/10.3389/fnhum.2022.861517

Fernandez-Duque, D., Grossi, G., Thornton, I. M., & Neville, H. J. (2003). Representation of Change: Separate Electrophysiological Markers of Attention, Awareness, and Implicit Processing. Journal of Cognitive Neuroscience, 15(4), 491–507. https://doi.org/10.1162/089892903321662895

Förster, J., Koivisto, M., & Revonsuo, A. (2020). ERP and MEG correlates of visual consciousness: The second decade. Consciousness and Cognition, 80, 102917–102924. 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, 34(5), 1738–1747. https://doi.org/10.1523/JNEUROSCI.4403-13.2014

Gaillard, R., Dehaene, S., Adam, C., Clémenceau, S., Hasboun, D., Baulac, M., Cohen, L., & Naccache, L. (2009). Converging intracranial markers of conscious access. PLoS Biology, 7(3), 0472–0492. https://doi.org/10.1371/journal.pbio.1000061

Graziano, M. S. A., & Kastner, S. (2011). Awareness as a perceptual model of attention. Cognitive Neuroscience, 2(2), 125–127. https://doi.org/10.1080/17588928.2011.585237

Graziano, M. S. A., & Webb, T. W. (2015). The attention schema theory: A mechanistic account of subjective awareness. Frontiers in Psychology, 6, 500–500. https://doi.org/10.3389/fpsyg.2015.00500

Grootswagers, T., Wardle, S. G., & Carlson, T. A. (2017). Decoding Dynamic Brain Patterns from Evoked Responses: A Tutorial on Multivariate Pattern Analysis Applied to Time Series Neuroimaging Data. Journal of Cognitive Neuroscience, 29(4), 677–697. https://doi.org/10.1162/jocn_a_01068

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, 32(19), 4139-4149.e4. https://doi.org/10.1016/j.cub.2022.07.068

Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and Overlapping Representations of Faces and Objects in Ventral Temporal Cortex. Science (American Association for the Advancement of Science), 293(5539), 2425–2430. https://doi.org/10.1126/science.1063736

Helmholtz, H. von, & Southall, J. P. C. (James P. C. (1924). Helmholtz's treatise on physiological optics. Optical Society of America.

Herman, W. X., Smith, R. E., Kronemer, S. I., Watsky, R. E., Chen, W. C., Gober, L. M., Touloumes, G. J., Khosla, M., Raja, A., Horien, C. L., Morse, E. C., Botta, K. L., Hirsch, L. J., Alkawadri, R., Gerrard, J. L., Spencer, D. D., & Blumenfeld, H. (2019). A Switch and Wave of Neuronal Activity in the Cerebral Cortex During the First Second of Conscious Perception. Cerebral Cortex (New York, NY), 29(2), 461–474. https://doi.org/10.1093/cercor/bhx327

Hesse, J. K., & Tsao, D. Y. (2020). A new no-report paradigm reveals that face cells encode both consciously perceived and suppressed stimuli. eLife, 9, e58360. https://doi.org/10.7554/eLife.58360

Humphrey, N. K., & Weiskrantz, L. (1967). Vision in Monkeys after Removal of the Striate Cortex. Nature (London), 215(5101), 595–597. https://doi.org/10.1038/215595a0

Hutchinson, B., Dehaene, S., Slagter, H., & Pitts, M. (2025). A late meta-stable code of conscious access in the absence of report.

Hutchinson, B. T., Jack, B. N., Pammer, K., Canseco-Gonzalez, E., & Pitts, M. (2024). No electrophysiological evidence for semantic processing during inattentional blindness. NeuroImage, 299, 120799. https://doi.org/10.1016/j.neuroimage.2024.120799

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), 1535–1535. https://doi.org/10.1038/s41467-022-28897-2

King, J.-R., & Dehaene, S. (2014). Characterizing the dynamics of mental representations: The temporal generalization method. Trends in Cognitive Sciences, 18(4), 203–210. https://doi.org/10.1016/j.tics.2014.01.002

King, J.-R., Pescetelli, N., & Dehaene, S. (2016). Brain Mechanisms Underlying the Brief Maintenance of Seen and Unseen Sensory Information. Neuron, 92(5), 1122–1134. https://doi.org/10.1016/j.neuron.2016.10.051

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

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 (American Association for the Advancement of Science), 340(6130), 376–380. https://doi.org/10.1126/science.1232509

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 of London. Series B. Biological Sciences, 373(1755), 1–12.

Mack, A., & Rock, I. (1998). Inattentional blindness. (pp. xiv, 273). The MIT Press.

Mashour, G. A., Roelfsema, P., Changeux, J.-P., & Dehaene, S. (2020). Conscious Processing and the Global Neuronal Workspace Hypothesis. Neuron (Cambridge, Mass.), 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

Melloni, L., Mudrik, L., Pitts, M., & Koch, C. (2021). Making the hard problem of consciousness easier. Science, 372(6545), 911–912. https://doi.org/10.1126/science.abj3259

Nguyen, B. N., Chan, Y. M., Bode, S., & McKendrick, A. M. (2020). Orientation-dependency of perceptual surround suppression and orientation decoding of centre-surround stimuli are preserved with healthy ageing. Vision Research, 176, 72–79. Scopus. https://doi.org/10.1016/j.visres.2020.07.015

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

Panagiotaropoulos, T. I. (2024). An integrative view of the role of prefrontal cortex in consciousness. Neuron, 112(10), 1626–1641. https://doi.org/10.1016/j.neuron.2024.04.028

Pashler, H. (1994). Dual-task interference in simple tasks: Data and theory. Psychological Bulletin, 116(2), 220–244. Scopus. https://doi.org/10.1037/0033-2909.116.2.220

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, 1078–1078. 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 (Orlando, Fla.), 101, 337–350. https://doi.org/10.1016/j.neuroimage.2014.07.024

Ramkumar, P., Jas, M., Pannasch, S., Hari, R., & Parkkonen, L. (2013). Feature-specific information processing precedes concerted activation in human visual cortex. The Journal of Neuroscience, 33(18), 7691–7699. https://doi.org/10.1523/JNEUROSCI.3905-12.2013

Ramsøy, T. Z., & Overgaard, M. (2004). Introspection and subliminal perception. Phenomenology and the Cognitive Sciences, 3(1), 1–23. https://doi.org/10.1023/B:PHEN.0000041900.30172.e8

Salti, M., Monto, S., Charles, L., King, J.-R., Parkkonen, L., & Dehaene, S. (2015). Distinct cortical codes and temporal dynamics for conscious and unconscious percepts. eLife, 4, e05652. https://doi.org/10.7554/eLife.05652

Schalk, G., Kapeller, C., Guger, C., Ogawa, H., Hiroshima, S., Lafer-Sousa, R., Saygin, Z. M., Kamada, K., & Kanwisher, N. (2017). Facephenes and rainbows: Causal evidence for functional and anatomical specificity of face and color processing in the human brain. Proceedings of the National Academy of Sciences - PNAS, 114(46), 12285–12290. https://doi.org/10.1073/pnas.1713447114

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., Pitts, M., Bruchmann, M., & Straube, T. (2020). Differential Effects of Awareness and Task Relevance on Early and Late ERPs in a No-Report Visual Oddball Paradigm. The Journal of

Neuroscience, 40(14), 2906–2913. https://doi.org/10.1523/JNEUROSCI.2077-19.2020

Schurger, A., Sarigiannidis, I., Naccache, L., Sitt, J. D., & Dehaene, S. (2015). Cortical activity is more stable when sensory stimuli are consciously perceived. Proceedings of the National Academy of Sciences - PNAS, 112(16), E2083–E2092. https://doi.org/10.1073/pnas.1418730112

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

Sergent, C., & Dehaene, S. (2004). Is consciousness a gradual phenomenon? Evidence for an all-or-none bifurcation during the attentional blink. Psychological Science, 15(11), 720–728. Scopus. https://doi.org/10.1111/j.0956-7976.2004.00748.x

Seth, A. K., & Bayne, T. (2022). Theories of consciousness. Nature Reviews. Neuroscience, 23(7), 439–452. https://doi.org/10.1038/s41583-022-00587-4

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

Sternberg, S. (1969). The discovery of processing stages: Extensions of Donders' method. Acta Psychologica, 30, 276–315. https://doi.org/10.1016/0001-6918(69)90055-9

Treder, M. S. (2020a). MVPA-Light: A Classification and Regression Toolbox for Multi-Dimensional Data. Frontiers in Neuroscience, 14, 289–289. https://doi.org/10.3389/fnins.2020.00289

Treder, M. S. (2020b). MVPA-Light: A Classification and Regression Toolbox for Multi-Dimensional Data. Frontiers in Neuroscience, 14. https://doi.org/10.3389/fnins.2020.00289

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

Watson, A. B. (2017). QUEST: A general multidimensional bayesian adaptive psychometric method. Journal of Vision (Charlottesville, Va.), 17(3), 1–27. https://doi.org/10.1167/17.3.10

Zhu, Y., Li, C., Hendry, C., Glass, J., Canseco-Gonzalez, E., Pitts, M. A., & Dykstra, A. R. (2024). Isolating Neural Signatures of Conscious Speech Perception with a No-Report Sine-Wave Speech Paradigm. Journal of Neuroscience, 44(8). https://doi.org/10.1523/JNEUROSCI.0145-23.2023