THE UNIVERSITY OF QUEENSLAND

AUSTRALIA

Determining the Role of Neural Oscillations in Perception Using Model-Based Approaches

Henry Antony Beale
BA (Hon I)

0000-0003-4176-378

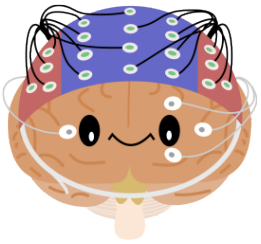

A thesis submitted for the degree of Doctor of Philosophy
at The University of Queensland in 2026.
Queensland Brain Institute

Abstract

A major goal in neuroscience is to understand how brain activity gives rise to cognition and perception. Posterior alpha (8–14 Hz) oscillations are consequential for visual perception (Jensen & Bonnefond, 2026), but it remains unclear how they affect the neuronal computations that support visual information processing. Research on the relationship between neural oscillations and perceptual experience has developed largely independently of major theoretical advances in computational neuroscience and psychophysics, including observer modelling and information-theoretic approaches to neural population coding. In my work, I combined behavioural testing with electroencephalography (EEG) and these computational approaches. The aim of this thesis is to advance current understanding of how alpha-band oscillations shape perception via their effects on sensory coding and neuronal computation.

In the first study of this thesis (Chapter 2), I examined how alpha oscillations impact the contrast response function, which describes how the strength of perceptual representations relates to external stimulus intensity. Human observers viewed stimuli across the full range of contrast intensities and performed both a detection and a discrimination task. I applied a Signal Detection Theory model that showed perceptual sensitivity was modulated by divisive suppression of neural response gain, affecting both the perceptual strength of stimuli and the baseline neural activity that contributes to false alarms in detection. This study is the first to model alpha power and phase as a combined inhibitory effect across the full range of psychometric performance. Although alpha signals are oscillatory and have been linked to pulsed inhibition of neural activity (Mathewson et al., 2011), I show that the resulting perceptual suppression is relatively sustained across successive cycles, with substantially weaker phasic modulation than proposed by pulsed-inhibition theories.

In Chapter 3, I examined how alpha oscillations alter neural representations of visual orientation. While participants performed an orientation estimation task, I used real-time EEG analysis to present stimuli during high or low power oscillations. During high alpha power, responses were more variable and biased further away from the horizontal and vertical orientation axes. Neural population codes produce repulsive estimation biases because they are adapted to the statistical distribution, or *prior probability*, of environmental information (Wei & Stocker, 2015) consistent with the efficient coding hypothesis (Barlow, 1961). Using a Bayesian observer model that allocates representational resources anisotropically

in line with environmental orientation statistics, I found that high alpha power increases sensory uncertainty, which is sufficient to explain stronger repulsive biases, and also increases guessing rate. These findings suggest that alpha power modulates sensory noise within efficient neural codes for orientation.

In Chapter 4, I developed an observer model that uses a spiking neural network to investigate the mechanistic links between alpha-frequency neural inhibition and visual information processing. The model provides a compact, theory-driven description of functioning visual neurons and tests predictions about how neural inhibition alters sensory coding and computation, extending upon phenomenological models used so far. The model was applied to the results of the previous chapters, and it captured the divisive suppression of baseline and stimulus-driven responses found for detection in Chapter 2, as well as the stronger alpha-related orientation biases found in Chapter 3. Model simulations suggested that alpha inhibition modulates neuronal gain through either reduced excitation or increased divisive normalisation arising from inhibitory connections within a neural population. Such inhibitory pathways are the basis of *surround modulation* in visual neurons, whereby receptive fields are modulated by activity in adjacent neurons. To examine this, Chapter 5 presents a behavioural experiment in which an oriented spatial surround stimulus was used to reset the phase of endogenous oscillations within a putative lateral inhibition network. By testing observers' ability to detect an oriented central stimulus at varying latencies, I found that perceptual sensitivity was modulated rhythmically depending on the relative orientation alignment of the surround and central stimulus, consistent with the orientation tuning of surround modulation circuits. This provides support for the involvement of lateral inhibition circuits in oscillatory visual processing.

Taken together, the results of my experimental and theoretical work suggest that alpha oscillations alter the fidelity of sensory representations conveyed by spiking activity in the visual cortex. This suppressive effect is proposed to be mediated primarily through firing-rate inhibition, and I show that its perceptual consequences depend strongly on the organisation of sensory codes and on the demands of specific perceptual tasks such as discrimination, estimation, and detection. These results point to a key role of alpha oscillations in organising visual sensitivity into a temporal code.

Declaration by Author

This thesis is composed of my original work, and contains no material previously published or written by another person, or generated by artificial intelligence (AI), except where clearly referenced and acknowledged either in the main text or the preliminary pages of this thesis. I have clearly stated the contribution of others to jointly authored works that I have included in my thesis.

I have clearly stated the contribution of others to my thesis as a whole, including statistical assistance, survey design, data analysis, significant technical procedures, professional editorial advice, use of AI tools, financial support and any other original research work used or reported in my thesis. The content of my thesis is the result of work I have carried out since the commencement of my higher degree by research candidature and does not include a substantial part of work that has been submitted to qualify for the award of any other degree or diploma in any university or other tertiary institution. I have clearly stated which parts of my thesis, if any, have been submitted to qualify for another award.

I acknowledge that an electronic copy of my thesis must be lodged with the University Library and, subject to the policy and procedures of The University of Queensland, the thesis be made available for research and study in accordance with the Copyright Act 1968 unless a period of embargo has been approved by the Dean of the Graduate School.

I acknowledge that copyright of all material contained in my thesis resides with the copyright holder(s) of that material. Where appropriate I have obtained copyright permission from the copyright holder to reproduce material in this thesis and have sought permission from co-authors for any jointly authored works included in the thesis.

Publications Included in this Thesis

Submitted

Beale. H., Mattingley, J. B., & Harris, A. M. (Under revision at *Nature Communications*) A joint alpha power–phase dynamic shapes visual perception. *bioRxiv*, 2026.02.19.706926. https://doi.org/10.64898/2026.02.19.706926

Drafted

Beale. H., Mattingley, J. B., & Harris, A. M. Alpha-band oscillations modulate sensory noise in efficient visual codes.

Beale. H., Gastrell, T., Williams., J. G., & Harris, A. Rhythmic visual perception arises from a surround modulation mechanism.

Other Publications During Candidature

Published

Harris, A. M., & **Beale. H.** (2025). Detecting behavioural oscillations with increased sensitivity: A modification of Brookshire's (2022) AR-surrogate method. *eLife*, 14:RP106141. https://doi.org/10.7554/eLife.106141.1

Nestor, P. J., ... **Beale. H.**, ... Gotz, J. (2025). A pilot safety and tolerability study of scanning ultrasound as a neuromodulation therapy in Alzheimer's disease. *Brain Communications*, 7 (6). https://doi.org/10.1093/braincomms/fcaf445

Williams, J. G., Harrison, W. J., **Beale. H.**, Mattingley, J. B., & Harris, A. M. (2024). Effects of neural oscillation power and phase on discrimination performance in a visual tilt illusion. *Current Biology*, 34 (8). https://doi.org/10.1016/j.cub.2024.03.014

Drafted

Beale. H. & Harrison, W. J. On the inseparabiliy of the prior and neural resources in behavioural bias. *bioRxiv*, 2026.04.07.714659. https://doi.org/10.64898/2026.04.07.7

Contributions by Others to this Thesis

The following work has been included as part of Chapter 2:

Beale. H., Mattingley, J. B., & Harris, A. M. (Submitted) A joint alpha power–phase dynamic shapes visual perception. *bioRxiv*, 2026.02.19.706926. https://doi.org/10.64898/2026.02.19.706926.

	Conceptualisation	**Data Collection**	**Analysis**	**Writing**
HB	70%	100%	100%	80%
JM	0%	0%	0%	5%
AH	30%	0%	0%	15%

The following work has been included as part of Chapter 3:

Beale. H., Mattingley, J. B., & Harris, A. M. (Drafted) Alpha-band oscillations modulate sensory noise in efficient visual codes.

	Conceptualisation	**Data Collection**	**Analysis**	**Writing**
HB	100%	100%	100%	90%
JM	0%	0%	0%	5%
AH	0%	0%	0%	5%

The following work has been included as part of Chapter 5:

Beale. H., Gastrell, T., Williams., J. G., & Harris, A. (Drafted) Rhythmic visual perception arises from a surround modulation mechanism.

	Conceptualisation	**Data Collection**	**Analysis**	**Writing**
HB	100%	90%	100%	90%
TG	0%	5%	0%	5%
JW	0%	5%	0%	0%
AH	0%	0%	0%	5%

Use of Artificial Intelligence

No artificial intelligence used.

Editorial Assistance

No editorial assistance.

Financial Support

This research was supported by an Australian Government Research Training Program Scholarship [DOI: https://doi.org/10.82133/C42F-K220] and by an Australian Research Council Discovery Early Career Researcher Award (DE220101019) under Dr. Anthony Harris.

Inclusion of Work Submitted for Another Degree

No works submitted towards another degree or diploma have been included in this thesis.

Research Involving Human or Animal Subjects

This research was approved by the University of Queensland Human Research Ethics Committee A under the project title *Understanding how neural oscillatory phase affects perception and attention* (2021/HE002284). A copy of the ethics approval letter is included in **Appendix A**.

Acknowledgements

The work in this thesis is the result of a wonderful collaboration that I have had with Dr. Anthony Harris over the last five years. I am grateful for his support and generosity. Thank you for being a kind man and a true gentleman. Thank you, Ant.

I would also like to thank Prof. Jason Mattingley for his support over the years that I have been in the lab. He has been a clear mind and a source of wisdom that I trust deeply.

Will Harrison (a great man who will not ever read this) is a great friend, wonderful cyclist (at some point in my PhD when we cycled together), and a very generous person who has shared a lot of knowledge about computational neuroscience that is important to me.

I am grateful for the support and friendship of Dr. Emily A-Izzeddin—who was my first friend in science and the person who made me feel comfortable in the lab!

My friends throughout my PhD were fantastic—Sylvie Loneragan, Caleb Stone, Mimi Dang, Zak Buhmann, David Palmer!

Thanks also to the many friends I have met because of science—Tim Cottier, Matt O'Donohue, and Benjamin Lowe.

I would also like to express my gratitude for Laura Wang, Andrew McKay, Jessica Williams, Ben Wang, Morgan McIntyre, Zac Wang, and Ben Carman.

Tim Gastrell and Isabelle Paine were my weekend support, and I am very thankful for their love and friendship.

Thank you Danni for your help with selecting the font of this thesis (and being a great pinball partner and beyond)!

For M and D

Contents

List of Figures

List of Tables

List of Abbreviations and Symbols

Abbreviations

AIC	Akaike Information Criterion
BDT	Bayesian Decision Theory
BF	Bayes Factor
BIC	Bayesian Information Criterion
BSEM	Baseline Sensory Excitability Model
CI	Confidence Interval
EEG	Electroencephalography
FI	Fisher Information
GABA(ergic)	Gamma-Aminobutyric Acid
GLM	General Linear Model
HDI	Highest Density Interval
ICA	Independent Component Analysis
LFP	Local Field Potential
LGN	Lateral Geniculate Nucleus (of the thalamus)
LN	Linear-Nonlinear
MEG	Magnetoencephalography
SDT	Signal Detection Theory
TMS	Transcranial Magnetic Stimulation
V1	Visual Area 1 (also known as primary visual cortex)
V4	Visual Area 4

Symbols

cpd	cycles per degree
dva	degrees of visual angle
Hz	Hertz (cycles per second)
ms	milliseconds
sd	standard deviation

Chapter 1

General Introduction

> For indeed it has been well said,
>
> 'In every object there is inexhaustible meaning; the eye sees in it what the eye brings the means of seeing.'
>
> To Newton and to Newton's dog Diamond, what a different pair of Universes; while the painting on the optical retina of both was, most likely, the same!
>
> Thomas Carlyle,
> *The French Revolution: A History*

> Rhythm is the most important thing in your goddamn life.
>
> Keith Richards,
> *The New York Times*

When you open your eyes the world floods in and dances among the activity of billions of neurons inside your brain. How these patterns of neural activity give rise to our experience of the world is a mystery at the heart of sensory and cognitive neuroscience research. Recordings of the brain reveal that sensory neurons dance *rhythmically*, around ten times per second, which is surprising given the seemingly uninterrupted experience of the world that we experience.

The study of the brain has revealed much about the anatomical and functional organisation of our cognitive and sensory systems. However, most of this work

has assumed a model of neural processes as a steady and feedforward flow of information in the brain. The presence of brain rhythms, on the other hand, implies that neural activity is unavoidably temporal (Kusnir & Landau, 2025). The overarching aim of this thesis is to understand how visual perception is affected by the time-varying nature of brain function.

1.1 A Brief History of the Alpha Rhythm

In 1924, the German psychiatrist Hans Berger made the first successful recordings of electrical signals from a human brain. Berger chose the name *Electroencephalogram*[1] for these recordings and noted in them the presence of large amplitude waves with a frequency of roughly 10 Hz, which he termed the *alpha* rhythm. These alpha waves were of particular interest because of their apparent modulation by sensory input and mental activity. Berger described these findings in a later publication (Berger, 1929) that documented his years-long experimentation and attempts to rule out the possibility that these measurements originated outside the brain, such as from a muscular or vascular source.

More than 100 years have passed since these original observations, yet remarkably the alpha rhythm remains the only signature of brain activity that can be read using the naked eye. In today's research laboratories, healthy volunteers can observe their alpha rhythms displayed in real-time on a computer screen, from an electroencephalogram recorded at the scalp. A common way to induce the alpha rhythm is to instruct the participant (as I have done for almost all of my volunteers) to close their eyes and then reopen them. Typically, a strong oscillation in the signal would occur with a period of ~100 milliseconds. These demonstrations are not far removed from those that were made originally by Berger! To this day, they serve as compelling evidence for the relevance of the alpha rhythm to sensory processing[2].

Researchers have been interested in whether alpha oscillations might play a role in sensory processing since the early days of EEG research. However, the neural origins of the alpha rhythm were not immediately accepted. Following their publication, Berger's results were initially treated with scepticism. In the 1930s,

[1] *Elektrenkephalogramm* in original German

[2] From a practical standpoint, they also confirm the EEG equipment is measuring signal from the brain and not simply displaying noise. Beyond that, many healthy people may go their entire lives without ever observing the direct activity of their brain, making this a fascinating and memorable moment.

two English researchers attempted to closely reproduce Berger's findings, sharing similar concerns around the potential non-cerebral origin of Berger's rhythms. After a series of control experiments, Adrian and Matthews (1934) hesitantly concluded that the observed electrical potentials do, in fact, arise from the occipital cortex (although they disagreed with Berger's claim that rhythms occurred widely across the brain).

An illustrative example of the recordings made by Adrian and Matthews is shown in Figure 1.1a. A key finding they reproduced is that the alpha rhythm appears quickly over the occipital cortex within half a second of the eyes being shut. In addition, they demonstrated the dependence of the rhythm on the absence of attentive visual processing, which is distinct from the absence of visual stimulation alone. Specifically, a participant whose head was enclosed in a black velvet-lined box with their eyes open would show the rhythm eventually after some time. Despite their eyes being open in the dark all along, giving the instruction to once more try to 'see' would result in these spontaneous rhythms being quickly extinguished. This provided a clear demonstration that the alpha rhythm was abolished during effortful visual processing. Adrian and Matthews (1934) summarised their findings as a tendency for visual neurons to "beat synchronously when they are undisturbed" (p. 356).

Early theories viewed the alpha rhythm through the lens of fluctuating excitability (Jasper, 1936), proposing that alpha might reflect a 'scanning' function (Walter & Walter, 1949) or a mechanism for sensory timing (Ellingson, 1956; Lindsley, 1952). These theories suggest that the nature of perception *is rhythmic*, with sensory content processed in discrete packages every ~100 ms (VanRullen, 2018). Multiple early experiments demonstrated that the phase of alpha oscillations predicted reaction times during visual tasks (Callaway & Yeager, 1960; Dustman & Beck, 1965; Lansing, 1957). Although this evidence implied that alpha rhythms played a role in sensory processing, reaction times reflect the culmination of many sensory and motor processes, making a mechanistic interpretation difficult. Additionally, other studies failed to find an association between alpha oscillations and behaviour (e.g. Walsh, 1952).

Nunn and Osselton (1974) showed some of the first general evidence that visual perception was phase-locked to alpha cycles (measured by galvanic skin responses to masked presentations of the word 'DANGER'). Another study by Milstein (1974) showed that brief visual stimuli reduced the alpha rhythm depending on the alpha phase at the time of stimulus presentation, suggesting that phase

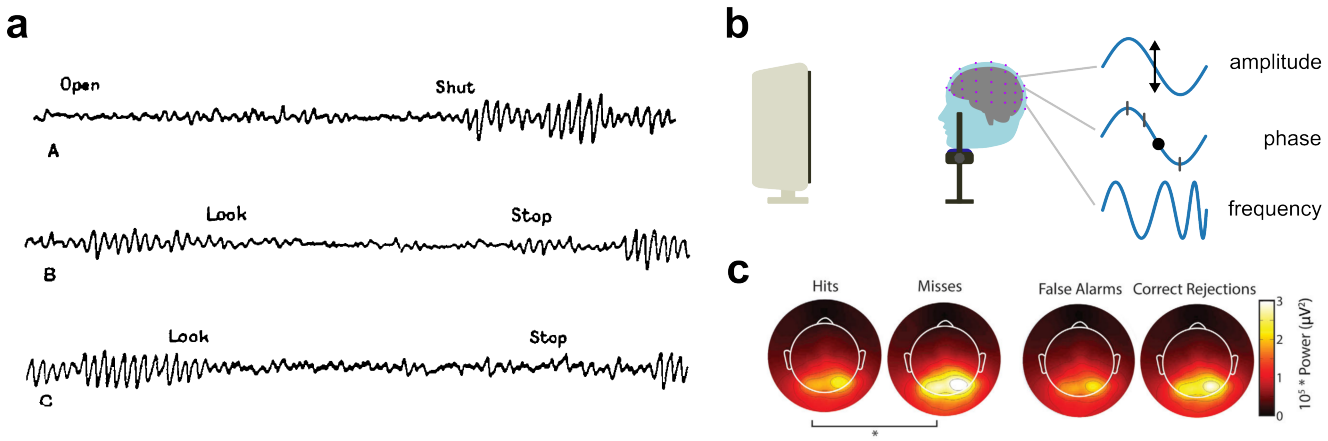

Figure 1.1: From velvet-lined box to liquid crystal display: Early and modern investigations on visual alpha rhythms

a Original recordings from Adrian and Matthews (1934) showing the occipital electroencephalogram of a person whose head is enclosed in a box lined with black velvet. Top trace: shutting of the participant's eyes produced a strong amplitude alpha rhythm. Middle and lower traces: two recordings that show alpha rhythms that emerged after several minutes in the dark with eyes open; the instruction to look (i.e. self-directed visual attention) abolished the rhythms. **b** Icons showing a modern recording setting with multichannel EEG during a computer-based visual task. Blue waveforms represent alpha waves which can be decomposed into their amplitude, phase, and frequency. **c** Scalp topographies of prestimulus alpha power from a 64-channel EEG recording during a visual detection task (Limbach & Corballis, 2016); sorted by response outcome (hit, miss, false alarm, and correct rejection). Asterisk denotes the significant increase in alpha power prior to the presentation of undetected (missed) stimuli, relative to stimuli that were seen by the participant. Figures from reprinted with the publishers' permission from Adrian and Matthews (1934) and Limbach and Corballis (2016).

might modulate the strength of the neural response to the stimulus. Following this, an influential paper by Varela et al. (1981) demonstrated that the alpha rhythm phase modulated the perception of apparent motion and judgements of simultaneity. To do this, they developed a closed-loop EEG system that allowed them to analyse brain activity in real time and trigger a light-emitting diode based on alpha oscillations measured on the scalp (see also Gho and Varela, 1988).

More recent work has continued the examination of how alpha power and phase might play a role in visual perception tasks (e.g. Busch et al., 2009; Iemi et al., 2022; Thut et al., 2006; discussed in greater depth later in this chapter; see also Figure 1.1b). Recent research has also taken advantage of standardised EEG equipment that records many electrodes over the scalp. An illustration of this is shown in Figure 1.1c from a study by Limbach and Corballis (2016). The topographic maps in Figure 1.1c show the distribution of alpha power over the scalp in a visual detection task. Limbach and Corballis (2016) found an increase in power over

occipital sensors before the presentation of stimuli that participants subsequently missed, consistent with the initial findings that the alpha rhythm is strongest during visual inactivity (Adrian & Matthews, 1934; Berger, 1929). However, the results of this study also showed increased alpha power prior to stimuli that were never presented and were falsely reported. This raises the question of whether alpha reflects observers' sensory processing of the stimulus or their propensity to respond with a particular decision. I will discuss this question in greater depth, below, in the section *Experimental Evidence for Rhythmic Perception*.

1.2 Where Does the Alpha Rhythm Come From?

Before discussing how alpha rhythms shape neural processing and perception, it is worth considering first how these rhythms are generated in the brain. Importantly, there may be more than one alpha generator. There is evidence of alpha oscillations in many brain regions, which vary in their specific frequency both within and across persons (Haegens et al., 2014, 2015; Meij et al., 2016). Recordings of local field potentials (LFPs) in monkeys show that oscillatory peak frequencies tend to increase in downstream regions of the cortical hierarchy, even when looking within commonly defined bands (i.e. *theta*: 2–8 Hz, *alpha*: 8–12 Hz, *beta*: 12–35 Hz; Lundqvist et al., 2020). Such variability suggests that the mechanisms underlying alpha rhythms may be heterogeneous throughout the brain. Despite this, some important principles and specific alpha-generating circuits are known, and these are briefly reviewed here.

The biophysical origin of the visual EEG alpha rhythm is likely to be pyramidal neurons in the occipital cortex (Lopes da Silva & Storm van Leeuwen, 1977; Silva et al., 1991). Pyramidal cells are a type of excitatory neuron that are spatially elongated and extend vertically through multiple layers of cortex. Their large spatial extent makes the transmembrane currents caused by their synaptic activity a significant contributor to the electrical fields that are measured by EEG (see Buzsáki et al., 2012, for review). Isolated pyramidal cells from layer V of the occipital cortex spontaneously generate 5–12 Hz rhythmic firing patterns due to their intrinsic membrane properties (Silva et al., 1991). However, single cells contribute only weakly to electric fields, meaning that the already weak signals measured by scalp EEG must reflect the *synchronised* activity of many thousands to millions of pyramidal cells that are aligned in parallel (Lopes da Silva, 2013).

How then does synchrony arise within the activity of neural populations? A compact computational explanation comes from a mathematical model developed

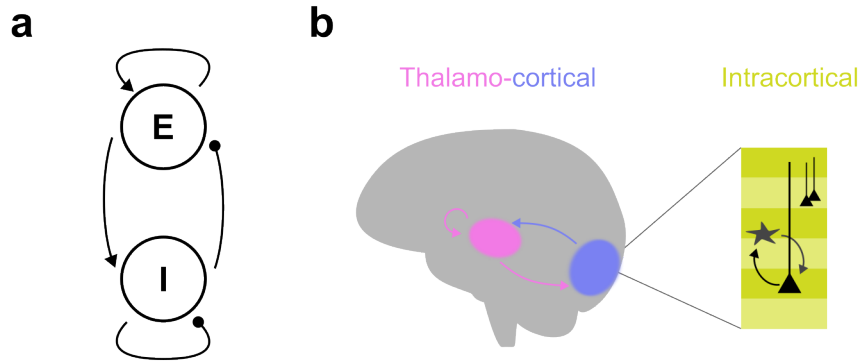

Figure 1.2: Circuit-level origins of the alpha rhythm

a A circuit diagram of recurrent connections between populations of excitatory (E) and inhibitory (I) neurons in the Wilson-Cowan model (Wilson & Cowan, 1972). Arrows show directed excitatory connections, whereas circles show directed inhibitory connections. **b** Two major generators of alpha oscillations in the visual system are neurons within thalamo-cortical pathways and the interlaminar intracortical circuitry of the cortex (summarised in the main text).

by Wilson and Cowan (1972) who studied the dynamics of interconnected neurons. They analysed the activations within subpopulations of simple excitatory and inhibitory neurons, which could be coupled to each other or recurrently self-connected (shown diagrammatically in Figure 1.2). Under certain parameter settings, the neural population exhibits sustained oscillations. Intuitively, this emerges from the *push-pull* relationship between excitation and lagged inhibition (i.e. excitation wanes as inhibition grows, causing a release from inhibition that promotes excitation and grows inhibition, and so forth). These computational findings highlight that oscillations emerge as a *network* property of the brain because of coupled interactions between excitatory and inhibitory neurons, occurring within and across brain areas.

A major source of occipital alpha oscillations is likely to involve recurrent connections between the thalamus and the primary visual cortex (V1). This thalamo-cortical loop (illustrated in Figure 1.2b) provides the neurophysiological mechanism by which visual signals can be transmitted to the cortex at an alpha frequency. In the visual system, signals from the retina are transmitted to the lateral geniculate nucleus (LGN) of the thalamus, a subcortical structure that has excitatory feedforward projections targeting the pyramidal cells in the middle layer 4 of V1 (Briggs, 2017). Pyramidal cells in layers 5/6 of V1 also provide feedback projections to subcortical areas, including the LGN (Briggs, 2020).

Specialised neurons in the LGN exhibit intrinsic bursting at an alpha frequency consistent with a 'pacemaker' role (Hughes et al., 2011; Lopes da Silva,

1991). Other LGN neurons exhibit 'relay-mode' spikes that directly transmit retinal signals to the cortex, and this activity is suppressed by GABAergic interneurons in the LGN that fire based on the phase of alpha pacemaker cells (Lőrincz et al., 2009). The result is feedforward excitation of the visual cortex at the alpha frequency that could generate the alpha activity observed in cortical neurons (Lopes da Silva, 1991). In support of this, chemical inhibition that targets LGN pacemaker cells suppresses alpha activity in both the LGN and cortex (Hughes et al., 2011). Importantly, feedback connections from deeper layers of V1 project back to relay neurons and interneurons in the thalamus, and this appears to play an important role in organising the LGN alpha activity (Bastiaens et al., 2025; Halgren et al., 2019; Jones, 2009).

In concert with the thalamo-cortical source, there is evidence that alpha is generated within specific layers of cortex (Bollimunta et al., 2008; Buffalo et al., 2011). Using laminar probe recordings in macaque V1, current-source density analyses trace the source of alpha to the deeper layers of the cortex (Spaak et al., 2012), although other analyses suggest that alpha sources may exist across many layers (Haegens et al., 2015). More generally, layer-specific oscillations may form a canonical cortical microcircuit that mediates feedback and feedforward processing within and between regions (Bastos et al., 2015; Mendoza-Halliday et al., 2024).

Alpha activity in deeper layers is thought to correspond with feedback activity whereas gamma activity in superficial layers characterises feedforward activity. A study by van Kerkoerle et al. (2014) showed that V4 microstimulation elicited alpha oscillations in V1, consistent with V4 feedback projections that terminate in layer 5 of V1. In contrast, stimulation of V1 resulted in gamma activity within V4, consistent with feedforward activity. Recordings from human intracranial EEG support the existence of this spectro-spatial motif and suggest that alpha activity propagates backward along the visual hierarchy as travelling waves (Halgren et al., 2019). Moreover, cortical alpha was found to lead and (Granger causally) predict thalamic alpha (Halgren et al., 2019). This suggests that the thalamo-cortical alpha mechanisms discussed above could result from feedback signals initiated by the cortex.

In summary, there is evidence that alpha oscillations originate from both thalamic and cortical mechanisms, which are functionally intertwined. The strong alpha rhythm measured using EEG primarily reflects the activity of pyramidal cells in the visual cortex, which are the cells that encode specific features in the visual environment (reviewed further below). Moreoever, oscillations arise from circuits that feature inhibitory interneurons and play a broader role in layer-specific

circuits that interact with feedforward gamma-band activity. This suggests that oscillations may play a key role in sculting aspects of perception.

1.3 How Does Alpha Impact Neural Processes?

It is clear from what I have presented so far that inhibition is a major aspect of alpha rhythms. Occipital alpha rhythms arise from inhibitory neuronal interactions and their functional relationship with psychological processes appears to be suppressive. In line with this, current theories suggest that alpha rhythms reflect a pulsed-inhibition of neuronal activity that shapes perceptual processing (Mathewson et al., 2011; Mazaheri & Jensen, 2010). In this section, I discuss neurophysiological evidence that supports pulsed-inhibition and its significance in more general neural function.

A key finding is that alpha appears to coordinate cortical spiking activity. Neural spiking activity can be measured alongside oscillations in the local field potential (LFP), often using non-human primates with electrodes placed onto the cortical surface. Many studies show that spiking activity and high-frequency gamma oscillations (a functional correlate of spiking) are modulated according to the phase of alpha cycles in the LFP (Bollimunta et al., 2008; Bollimunta et al., 2011; Buffalo et al., 2011; Dougherty et al., 2017; Spaak et al., 2012; van Kerkoerle et al., 2014). Specifically, firing activity is greatest at the trough of alpha oscillations and weakest at the peak (e.g. Haegens et al., 2011). Furthermore, the strength of neuronal firing correlates inversely with the magnitude of alpha oscillations (Haegens et al., 2011; Iemi et al., 2022; van Kerkoerle et al., 2017).

Given that increasing alpha power appears to reduce the mean level of spiking activity, this implies that alpha's effects are primarily inhibitory. In other words, the oscillatory modulation of spiking activity is asymmetric. That is, this suggests that neural activity is gradually inhibited and then released from inhibition across the alpha cycle—without any phases that would appear to boost spiking activity beyond its uninhibited mean level (which are seen when alpha power is absent, for instance). These effects suggest spiking depends on alpha power and phase jointly (summarised in Figure 1.3). An asymmetric profile of inhibition could explain why increased alpha power is associated with reduced detection performance despite averaging over multiple phases (Mathewson et al., 2011). Analyses of alpha rhythms recorded from scalp EEG confirm that their waveforms are asymmetric and show that this property underlies late stimulus-evoked responses (Iemi et al.,

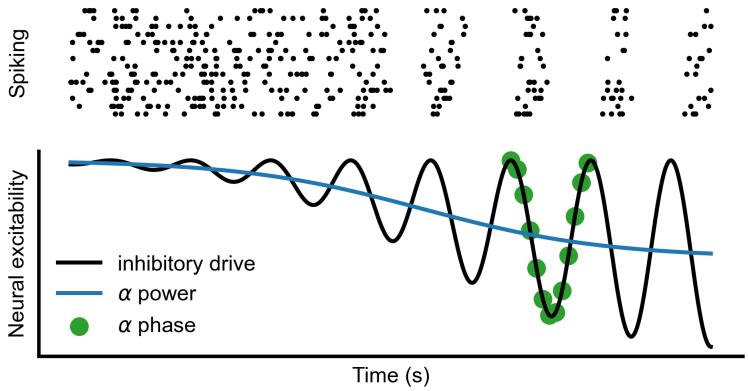

Figure 1.3: Pulsed inhibition of spiking activity

Rhythmic suppression of neural spiking activity at ~10 Hz (upper panel: raster plot showing the firing of a population of simulated neurons). The plot shows rhythmic inhibition emerging gradually toward the right. Stronger pulses of inhibition spaced 100 ms apart result in synchronous firing. Lower panel shows inhibitory pulses (black line) in population level activity reflected in local field potentials or scalp EEG. Pulsed-inhibition underlies measures of alpha-band activity such as power, which increases as a consequence of stronger inhibitory pulses (blue line), and phase (green dots).

2019; Mazaheri & Jensen, 2008; Nikulin et al., 2007). In general, the evidence supporting the pulsed-inhibition account suggests that alpha rhythms promote cyclic periods of suppressed neural excitation.

This raises the question: what are the impacts of pulsed-inhibition on neural processes and computations? A natural first-order consequence is the periodic gating of stimulus information that is conveyed in spiking activity. When alpha power is high, a functional consequence of this may be to promote the suppression of task-irrelevant information within neural networks that represent competing information in distinct cortical subregions (Jensen, 2024; Jensen & Mazaheri, 2010; Jensen et al., 2012). Pulsed-inhibition could also provide frequency-specific gain control that reflects reduced processing of inputs into a cortical circuit (Womelsdorf et al., 2014). Adjusting the gain of neural responses, such as via the activity of GABAergic interneurons, can modulate neural sensitivity to inputs without altering underlying selectivity (Ferguson & Cardin, 2020).

Beyond the impacts of pulsed-inhibition in local cortical circuits, a dominant view is that alpha promotes flexible neuronal communication between brain regions. Spatially and functionally distinct brains regions could become functionally linked, or segregated, based on the alignment of alpha phases across regions (Chapeton et al., 2019; Fries, 2005, 2015). This offers a way for distant units to

selectively communicate within the larger cortical networks relevant for perception and cognition, based on the synchrony of local inhibitory gating (Bonnefond et al., 2017). Such a mechanism would be ideal for attentionally guided processing of complex visual scenes that present the brain with competing stimuli/objects (Bonnefond et al., 2024).

Oscillations can also have computational benefits for the efficiency of neural processes. Biological systems operate under energy constraints, and this suggests that the brain should exploit neural processes that achieve its algorithmic goals with minimal metabolic expenditure. One suggestion is that oscillations are an efficient way to maintain encoded information in the presence of synaptic timing delays. This is because rhythmic bursts, as opposed to independent random spiking, require fewer neurons to fire and are less susceptible to representational drifts caused by noise (Chalk et al., 2016). Rhythmic processing allows neural representations to be distributed over time instead of over spatial structure. For instance, a phase-based code could segment and prioritise distinct representations, such as stimulus representations that fire preferentially at separate alpha phases based on pulsed-inhibition (Gips et al., 2016; Jensen et al., 2014; Masquelier et al., 2009). Additionally, oscillations can be 'multiplexed', allowing separate streams of activity to be distributed across frequency-bands but carried within the same spiking activity without interference (Akam & Kullmann, 2010, 2014). This could maximise the bandwidth of information embedded in local spiking activity, allowing a single channel to have distributed representations that can be flexibly rerouted.

Overall, alpha rhythms appear to result in pulsed-inhibition of neural spiking activity. At the local circuit level, this is likely to result in cycles of inhibitory gating and gain modulation. More broadly, these local oscillations might allow flexible processing of information between brain regions and could promote efficiency in distributed neural representations. Beyond this, it is clear that the precise effects of pulsed-inhibition will depend greatly on the context of the local circuits under consideration. This suggests the precise way in which alpha rhythms shape visual perception should depend on the visual information that is conveyed in spiking activity.

1.4 Linking Alpha Rhythms with Perception

Experimental Evidence for Rhythmic Visual Perception

The occipital alpha rhythm suggests some aspects of visual perception may be rhythmic due to the phasic fluctuations in the activity of neurons in the visual pathway. It is worth noting, however, that evidence for rhythmic perception that does not rely solely on measuring brain activity. Early studies analysed the distributions of reaction times in response to visual stimuli and found evidence for response modulation at a period consistent with alpha (Callaway & Alexander, 1960; Venables, 1960). More recently, behavioural experiments have shown that visual performance (e.g. detection and discrimination) fluctuates at an alpha frequency (Fiebelkorn et al., 2013; Landau & Fries, 2012; Michel et al., 2021; Re et al., 2019, 2023; Song et al., 2014). Furthermore, alpha rhythms in visual discrimination and contrast sensitivity have been found that are synchronised with the onset of motor actions and eye movements (Benedetto & Morrone, 2019; Benedetto et al., 2016; Tomassini et al., 2015; Zhang et al., 2019).

Without electrophysiology, a methodological requirement of these studies is that performance is analysed relative to a temporal reference point. It is typically assumed that neural oscillations are 'phase-reset' by a visual stimulus or attentional event (such as the onset of the trial or an informative pre-cue). By presenting target stimuli at varying latencies relative to the reset stimulus, rhythmic behaviour can be seen in the Fourier transformation of performance across trials (known as a *dense sampling* method; Kienitz et al., 2021; Tosato et al., 2022). Thus, behavioural experiments suggest that visual perception may have inherent rhythms, consistent with the suggestion of rhythmic perception theories. It is clear, however, that to provide a complete understanding of how rhythmic brain activity interacts with perception, it is necessary to measure brain activity.

Two original studies that linked rhythmic perception to EEG alpha rhythms were published in 2009. Using different stimuli and analysis approaches, Busch et al. (2009) and Mathewson et al. (2009) presented brief (~10 ms) visual stimuli near the threshold of visibility and analysed the relationship between prestimulus oscillatory phase and the proportions of correctly detected vs missed stimuli. Both studies found a correlation between alpha phase and visual detection, albeit in frontal and central EEG sensors, as opposed to occipital regions. A close replication of Busch et al. (2009) was conducted using MEG, and it found the expected phase effects were source-localised to occipital-temporal cortical regions

(Zazio et al., 2021). Furthermore, the phase of occipital alpha was shown to modulate the detection of phosphenes (i.e. illusory perceptions) caused by the delivery of transcranial magnetic stimulation (Dugué et al., 2011), suggesting these oscillations correspond to fluctuations in the excitability of visual cortex.

The correlation between prestimulus alpha phase and visual detection is a cornerstone result in support of rhythmic visual perception. Moreover, phasic modulation of perception appears to be consistent with the pulsed-inhibition proposal of how alpha oscillations influence neural activity. Specifically, phase modulates detection most strongly during high alpha power, where pulsed-inhibition is suggested to be strongest (Alexander et al., 2020; Fakche et al., 2022; Mathewson et al., 2009). However, the effects of phase on perceptual processes are left unclear by the original studies because they were missing stimulus-absent trials that would be required to model perceptual processes that underlie these effects (recall from earlier in this chapter that these were featured in the study by Limbach and Corballis, 2016; Figure 1.1c).

For an example of this problem, consider how participants performing a detection task might alter their performance by strategically choosing to always (or never) say something is present. This means that detection responses could arise from a mixture of the sensory system's processing of visual stimuli and also the observer's decisional goals. This fact is accounted for in signal detection theory (SDT) models (Tanner & Swets, 1954). SDT provides a process model that breaks perceptual decisions into two factors: (1) the *sensitivity* of the neural system to visual signals and (2) a decision *criterion* (see Figure 1.4a). Specifically, SDT assumes that stimuli are encoded into a noisy sensory response that differs on each trial. On trials with a stimulus present the sensory response is larger in magnitude (due to it being driven by the visual feature). SDT then assumes a decision rule whereby observers simply respond 'stimulus-present' if the signal exceeds a criterion level, otherwise they report 'stimulus-absent'. In this way, accuracy in the task can reflect either changes to how sensitive the observer is to the external signals or changes in the placement of their decision criterion (i.e. response biases).

The question of alpha's effects on perceptual processing has been more extensively studied with respect to alpha power than alpha phase. Several early studies showed that visual detection is diminished when stimuli are presented during periods of high prestimulus alpha power (Ergenoglu et al., 2004; Hanslmayr et al., 2007; Romei et al., 2008; van Dijk et al., 2008), consistent with pulsed-inhibition. Using SDT analyses, this effect was attributed to a change in response-bias or

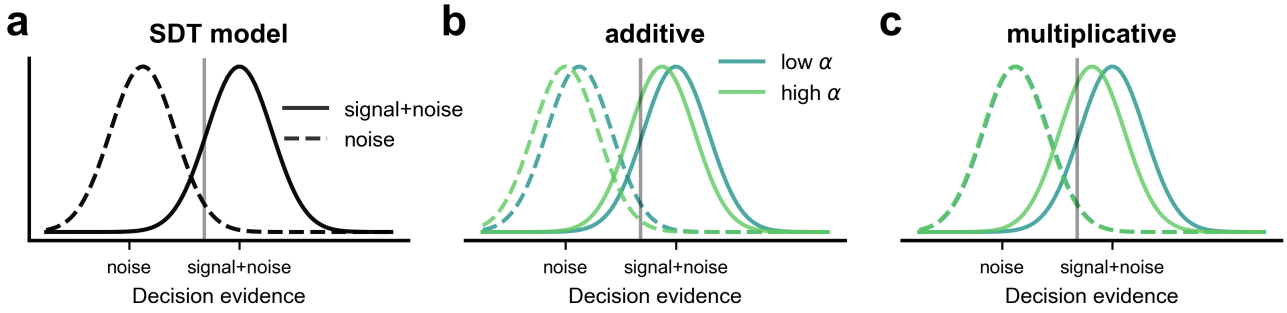

Figure 1.4: Signal detection theory model

a An overview of the signal detection theory model (Tanner & Swets, 1954) as applied to a simple detection task. Internal sensory responses are represented on the horizontal axis, and the Gaussian distribution reflects the variation of response magnitude across trials (due to sensory noise). On trials where a stimulus is presented, the sensory signal is shifted rightwards toward higher values (solid black line) but remains low when the stimulus is not presented (dashed line). The grey vertical line represents a decision criterion against which an observer might compare the sensory signal on a given trial and respond 'yes'/'stimulus-seen' if it exceeds the criterion. **b** Hypothetical effects of alpha oscillation modulation that produce an additive horizontal translation to sensory signals. Note how shifting the distributions of sensory signals is akin to moving the criterion in the opposite direction. **c** Hypothetical multiplicative modulation effects that depend on (or scale with) the strength of the sensory signal.

criterion rather than visual sensitivity (Iemi et al., 2017; Limbach & Corballis, 2016). While this suggests that alpha power promotes a conservative bias in decisions about stimulus presence, another account suggests that bias changes can be induced by changes in sensory cortical excitability (Samaha et al., 2020).

The baseline sensory excitability model (BSEM) was proposed to make sense of alpha's effects on visual detection biases, while also reconciling the fact that several alpha power studies do not find effects on visual discrimination accuracy (Samaha et al., 2020). The model proposes that activity in sensory cortices is elevated by the increased neural excitability that accompanies reduced alpha power. This could then influence detection without truly altering decision bias/criterion, because sensory noise provides an additive boost to decision signals that exceeds over the threshold for detection (Figure 1.4b). Thus sensory excitability could mimic the effects of moving the decision criterion in the opposite direction, becoming 'liberal' in high excitability states and 'conservative' in low excitability states. The influence of alpha power on the psychometric function of contrast seems to support an additive shift in base excitability (Pilipenko & Samaha, 2024), but other work has instead supported an alternative multiplicative gain modulation effect (Chaumon & Busch, 2014) whereby the modulation of sensory responses scales with stimulus intensity (Figure 1.4c).

It is certainly the case that alpha's excitability-based effects are more nuanced than implied by SDT's parcelling of decision-making into perceptual and decisional effects (Samaha et al., 2020). A recent study, showed that hemispheric lateralisation of alpha power predicts subjective contrast appearance for comparisons of stimuli in separate visual hemi-fields (Balestrieri & Busch, 2022). Interestingly, the presentation of longer duration stimuli (1.5 s) does not show an alpha power effect on criterion in detection as the BSEM predicts, and also shows effect on orientation discrimination counter to the evidence motivating the BSEM (Michail et al., 2021). A clever study by Zhou et al. (2021) used task instructions to manipulate participants' decision criteria in separate experiment blocks. This produced response bias changes without altering spontaneous prestimulus power, suggesting alpha may not reflect decisional factors. Moreover, the authors also found that inducing a conservative decision bias promotes a negative relationship between power and visual sensitivity, which suggests that low alpha facilitates neural information depending on participants' perceptual expectations. Broadly, the findings of Zhou et al. (2021) are consistent with the multiplicative gain modulation effects shown by Chaumon and Busch (2014), compared with the additive enhancements implied by the BSEM (Pilipenko & Samaha, 2024; Samaha et al., 2020).

A further challenge to understanding rhythmic perception is that it is unclear how the effects of alpha power and phase jointly modulate perceptual processes. While SDT analyses of alpha power in detection suggest a baseline/bias effect, a recent study has argued that alpha phase modulates noise in the SDT model (Pilipenko et al., 2026). Moreover, some studies find perceptual effects of alpha power but not phase (Benwell et al., 2017, 2022; Melcón et al., 2024). Disparate influences on perceptual processing by power and phase is a problem for the pulsed-inhibition model which posits a singular inhibitory impact on neural spiking. It is also concerning that some phase effects have failed replication attempts (Ruzzoli et al., 2019). Other lines of research have shown that alpha *frequency* modulates the temporal integration of brief, rapidly presented stimuli (Samaha & Postle, 2015; Wutz et al., 2018). These perceptual effects have similarly come under closer scrutiny following failed replication attempts (Buergers and Noppeney, 2022; Schoffelen et al., 2024; but see Samaha and Romei, 2024).

In summary, studies of alpha power and phase have provided results that are broadly consistent with the pulsed-inhibition hypothesis. However, it is currently unclear whether the perceptual modulations produced by changes in alpha phase

are consistent with those produced by changes in alpha power, as the theory would propose. It is also difficult to pinpoint the perceptual mechanisms that could be targeted by pulsed-inhibition. These processes are likely to feature inhibitory gain modulation, and a major question is whether this reflects an additive or multiplicative effect on the strength of visual representations (e.g. Figure 1.4b–c).

Replication failures also make the empirical status of rhythmic perception more uncertain. This is especially a concern for the supposed joint relationship between power and phase, as I have pointed out above. Mixed empirical findings may be a product of the different stimuli and experimental tasks used to investigate rhythmic perception, which are rarely replicated across studies. This variability could be a consequence of the field lacking an overarching theoretical understanding of how rhythmic brain activity influences (visual) perception, and a 'standard paradigm' to elicit these effects (as noted in a recent editorial; Keitel et al., 2022). A solution to this is to develop a model-based understanding of rhythmic perception and its neural underpinnings, which is informed by theoretical knowledge of visual processing.

This is, ultimately, the thesis that I will present.

A Brief Look at Neural Function in the Visual System

Here, I introduce some relevant theory of neural functions in the visual system. Note that the major ideas in this section form the basis of the coming thesis chapters and are revisited therein. Further, computational modelling of visual system neurons is central to Chapter 4, which contains a more rigorous mathematical treatment.

The purpose of the visual system is to process information about the world that is contained in patterns of light that enter the eye. A mosaic of receptors in the retinae of the eyes transduce light signals into neuronal impulses, using photochemical processes. At this point, visual information corresponds to the spatial patterns among adjacent photoreceptors (*rod* and *cone* cells) that will leave each retina via retinal ganglion cells that project to the lateral geniculate nucleus (LGN) of the thalamus, in each hemisphere (Briggs, 2017). There are three parallel processing streams in the LGN (parvo-, magno- and konio-cellular) that inherit receptive field properties from their afferent retinal ganglion cells. Importantly, these receptive fields are defined spatially as centre-surround and also by opponent-coded to light wavelength (red-green, blue-yellow, or luminance ON-OFF). This means that the fundamental content of the representation of early visual informa-

tion may be a *contrast code* to discriminate differences in the light environment (Adelson & Bergen, 1991).

The parallel streams of the LGN project to distinct layers in the primary visual cortex (area V1). Here, the neurons develop more advanced receptive fields that are selective for orientation, direction, color, and binocular disparity (Conway & Livingstone, 2006; Cumming & Parker, 1997; Gur et al., 2005; Hubel & Wiesel, 1968; Johnson et al., 2001) These receptive fields are, again, developed by the feedforward synaptic connections from afferent LGN neurons (Alonso et al., 2001). For instance, simple cells in V1 become orientation-selective by combining inputs from center-surround LGN receptive fields that are spatially aligned (Hubel & Wiesel, 1968). Similar feedforward connectivity rules can explain increasingly complex receptive fields in higher-level visual regions.

At each step, receptive fields enlarge and become more feature specific, segmenting into distinct ventral and dorsal streams that process 'what' vs 'where/how' information, respectively (Balaban & Ullman, 2025; Schneider, 1969). Perhaps the reason for this is that information becomes available for increasingly meaningful and behaviourally specialised purposes as signals progress up the visual hierarchy (Barlow, 1972). For example, neurons in dorsal stream area MT (V5) are selective to speed and direction (Adelson & Bergen, 1985; Rust et al., 2006), useful for interacting with moving and spatial features of the environment. In contrast, neurons in the ventral stream process the properties of objects, such as combinations of orientation, like angles/corners (V2; Anzai et al., 2007) and shape (V4; Pasupathy and Connor, 2002) that determine the recognition of objects and surface materials. Ultimately, at some point in this journey through the brain's receptive fields, higher areas like parietal/frontal cortex are no longer exclusively visual but multimodal or even 'amodal'. Thus, the visual system will have fulfilled what may be its primary task of "deliver[ing] a small set of useful measurements" about the light environment (Adelson & Bergen, 1991, p. 8) that feed into downstream processes which more directly govern behaviour.

There are three major computational functions of visual system neurons that are relevant to the current discussion: surround modulation, transducer functions, and population coding.

First, consider a problem faced by the visual system: inputs from the natural world are highly structured and highly redundant. Analysing the statistical structural of natural movies reveals that orientation content is correlated in space and time (Dong & Atick, 1995; van Hateren & Ruderman, 1998). This poses

a challenge because neurons, at a local level, would carry redundant meaning incurring an energetic cost from unnecessary spiking activity. In contrast, an *efficient coding* hypothesis (Barlow, 1961) suggests that a neural system evolved to run with minimal energetic expenditure should use a non-redundant code where each spike transmits maximum information about a stimulus (Olshausen & Field, 1996; Vinje & Gallant, 2000). A computational solution is inhibitory gain control, or *divisive normalisation*, to de-correlate redundant neural signals (Carandini & Heeger, 2012; Schwartz & Simoncelli, 2001). At a single neuron level, this corresponds to **surround modulation**, whereby stimulation outside a neuron's receptive field inhibits its responses. This property is found in neurons throughout sensory systems (reviewed in Angelucci et al., 2017). An example is the centre-surround spatial receptive fields of early visual neurons, in which neurons fire selectively in response to the *difference* between activity at the centre versus surround spatial region. Furthermore, neurons in higher-level visual regions show surround modulation for features like orientation (e.g. the tilt illusion) which may be an important part of how the brain processes spatiotemporal context in the environment (Schwartz et al., 2007).

A further way to understand the computations of visual neurons is by considering their **transducer functions**. A transducer function explains the functional relationship between external stimuli and their encoded representations. For simple one-dimensional stimulus features, this could refer to how orientation corresponds to a selectively tuned neural response (e.g. Hubel and Wiesel, 1962), or how light intensity maps to a level of spiking output (e.g. Naka and Rushton, 1966). Transducer functions are typically non-linear due to biophysical constraints of the nervous system and, potentially, efficient coding (Ganguli & Simoncelli, 2016). At a more abstract level, this shapes perception and determines how variations in the world affect psychological experience (Fechner, 1860). For example, the perceived intensity of many sensory attributes follows a power law function of the physical quantity (Stevens, 1957). Perception is also influenced by the noise variability of encoded representations, and the interaction of noise with the transducer function shapes the functional form of psychophysical thresholds (García-Pérez & Alcalá-Quintana, 2007; Zhou et al., 2024).

Finally, **population coding** reflects the fact that sensory variables evoke activity distributed across many neurons. Distributed representations are robust to noise variation and can protect encoded information from the failure of a small number of neurons (Pouget et al., 2000). Visual features can be encoded, in their full

range, by individual neurons that respond more selectively to local feature values. Thus, the precision of population-coded information is controlled by individual neurons' gain and selectivity (Brunel & Nadal, 1998). The efficient coding hypothesis posits that the predominant signals (e.g., the prevalence of vertical and horizontal information in our visual world) should receive greater representation from individual neurons in the population (Ganguli & Simoncelli, 2014). Accordingly, tuning preferences of visual neurons reflect a statistical *prior* distribution from the environment (Ganguli & Simoncelli, 2016); an asymmetric allocation of information that shapes perceptual behaviours like discrimination (Wei & Stocker, 2017). Separately, divisive normalisation and surround modulation can be seen as a transformation of the population response, allowing individual neurons to moderate their responsiveness based on the context of sensory activity (Heeger, 1992; Schwartz & Simoncelli, 2001). Therefore, population codes govern neural sensitivity to visual information. Downstream neurons in the brain can 'read' the population code, or integrate it into further computations, using simple feedforward synaptic connections (Pouget et al., 2003).

In summary, the visual system is the brain's primary interface with the visual world. Its neurons analyse sensory inputs and perform computations that determine perception and behaviour. I have highlighted three main ideas that characterise the function of visual neurons: surround modulation, transducer functions, and population coding. These functions matter because they describe the core computations carried out by sensory neurons, which will be crucial to understanding how oscillatory activity in these neurons influences perception.

What is Needed to Make Further Progress

Returning to the alpha rhythm, it is apparent that relatively little is understood about its relationship to the specific neural processes that govern perception. It is unclear how the putative neural basis of alpha, pulsed-inhibition, shapes visual processes like those mentioned in the previous section. Empirical studies that investigate alpha-linked changes in known visual computations—and move beyond characterising behavioural phenomenology—could greatly advance a theory-driven account of rhythmic perception.

Some key questions are: How do alpha rhythms shape the strength of perceptual representations in relation to known transducer functions of the visual system? Do alpha rhythms modulate the precision of perception or the noise properties of neural activity? What is the impact of pulsed-inhibition on popula-

tion codes? Does pulsed-inhibition share circuits that overlap with the inhibitory processes that operate in vision (e.g. surround modulation)? How can models of human behaviour incorporate realistic neural constraints, like alpha rhythms, to examine the effects those neural properties have on perception?

Psychophysical and model-based investigations provide a principled way to understand alpha's perceptual effects. This is facilitated by using stimuli designed to selectively target the activity of an established neural channel through which task-relevant information must flow. Many past studies that have linked alpha to visual processing have instead used arbitrary stimuli such as letters (Benwell et al., 2022), backward-masked shapes (Limbach and Corballis, 2016; Mathewson et al., 2009), or simple points of light (Busch et al., 2009; Zazio et al., 2021). Rhythmic perception should depend on how sensory stimuli are processed along the visual hierarchy and, crucially, whether oscillations intervene at relevant stages (Harris, 2023; Williams et al., 2024). The selection of behavioural tasks can also influence what computations must transpire in the brain to give rise to behaviour. For example, recall that the baseline sensory excitability model is motivated by contrasting results for visual detection versus discrimination. Ideally, a model-based account should link visual information processing to behaviour, and thus make alpha's effects in each task explicit.

A useful way to understand rhythmic perceptual processing is to use an *observer model*. Observer models attempt to map, quantitatively, the information processing steps that lead from stimulus to behaviour. Typically, they make explicit how stimulus information is encoded into an internal response that is subjected to a decision rule, while accounting for possible noise at each processing stage. Signal detection theory is a simple example of this (Tanner & Swets, 1954), but there are various other types of observer models of the visual system that have incorporated Bayesian inference (Yuille & Bülthoff, 1996), 'image-computable' inputs (e.g. Schütt and Wichmann, 2017, (Hermes et al., 2019)), or neural population coding (Jazayeri & Movshon, 2006; Ma et al., 2006). Observer models also provide a benchmark against which to test human performance (Geisler, 2011). While some alpha oscillation studies have used signal detection theory (Iemi et al., 2017; Limbach & Corballis, 2016; Zhou et al., 2021), more sophisticated observer models could clarify how pulsed-inhibition might impact neural encoding and information processing.

1.5 Overview of Research in this Thesis

The aim of this thesis is to uncover how alpha oscillations shape specific neural processes that underlie visual perception. In the following chapters, I present the results of three empirical studies and a further model-based investigation.

In Chapter 2, I present work that examined how alpha power and phase jointly affect the human contrast response function. I used a small-*N* EEG design to characterise the psychometric functions of observers in separate detection and discrimination tasks. A signal detection model that incorporates a contrast transducer function was used to explain psychometric data. Moreover, a novel link function was applied that estimates pulsed-inhibition state from the observed EEG phase and amplitude. Several models were tested that examine the influence of pulsed-inhibition on encoded visual responses. The results suggest that alpha oscillations relate to changes in response gain (concomitant with criterion in detection only) as opposed to other effects like threshold or noise influences. Interestingly, the influence of phase was found to be weaker than implied by existing accounts of pulsed-inhibition, suggesting that alpha may induce tonic inhibition that scales with alpha amplitude.

In Chapter 3, I examined whether alpha oscillations influence the perception of orientation and the repulsive orientation biases that result from efficient population coding. I used a real-time, closed-loop EEG processing design to present brief stimuli during periods of high or low alpha power, while participants performed an orientation reproduction judgement. The reported orientations were repulsively biased away from the cardinal horizontal and vertical axes that dominate the visual environment, consistent with efficient coding theory. Importantly, orientation biases were stronger for stimuli presented during periods of high alpha power. I used a Bayesian observer model to quantify anisotropic neural resource allocation and noise variability from behavioural reports. This revealed that alpha power increases the representational noise in efficient population codes, leading to biased perceptual estimation.

Chapter 4 presents a neural observer model to investigate the effects of alpha rhythms on visual processing. The model implements pulsed-inhibition in the form of spiking rate inhibition in a modelled network of orientation selective neurons. The model is fitted to the psychometric functions produced in Chapter 2 and shows computational support for the response gain modulations that were found. The predictions of the model are also consistent with the stronger orientation biases during high alpha found in Chapter 3. A novel computational

result was found that suggests that inhibition of the spiking rate could be mediated through either response gain modulation or divisive normalisation by lateral inhibitory connections.

In Chapter 5, I tested the predictions of the neural observer model in a behavioural experiment. Specifically, I examined whether the surround modulation circuitry in the visual cortex mediates rhythmic perception. A dense sampling paradigm was used to test detection performance over time after an oriented surround stimulus was used to phase-reset putative cortical rhythms. The results showed that visual sensitivity was modulated at an approximate 5 Hz rhythm, in a way that depended on the correspondence between the orientation of the reset stimulus and the target. These results are explained by a model simulation of oscillating neurons with lateral inhibitory connections.

The implications of these findings for a theoretical understanding of rhythmic visual perception are discussed in Chapter 6.

Chapter 2

A Joint Alpha Power–Phase Dynamic Shapes Human Visual Sensitivity

Human visual behaviour relies on inferences about the external world that are shaped by both sensory input and the internal states of the brain's sensory system. Even when sensory input is held constant, perceptual judgements can vary substantially owing to rhythmic oscillations in the excitability of sensory neural populations (Iemi et al., 2022). While it is well-accepted that neural oscillations are ubiquitous in sensory cortices and are linked to a variety of cognitive functions (e.g. visual attention, prediction, memory; Clayton et al., 2018; Pascucci et al., 2025), their specific impact on the computations underlying visual behaviour remains an open question.

Neural oscillations are a widespread feature of brain activity that are preserved across species (Buzsáki et al., 2013) and arise from neuronal micro-circuitry that features inhibitory dynamics (Womelsdorf et al., 2014). Low-frequency oscillations, particularly in the alpha band (8–14 Hz), reflect cycles of cortical excitability that shape the temporal dynamics of neural activity and putatively coordinate information processing amongst brain regions (Bonnefond et al., 2017; Chapeton et al., 2019; Fries, 2015; Jensen & Mazaheri, 2010). In sensory cortices, neuronal spiking is modulated by both the phase and amplitude of alpha oscillations (Dougherty et al., 2017; Haegens et al., 2011). Theoretical accounts suggest that alpha oscillations reflect a rhythmic pulsed-inhibition of neural activity (see Fig. 2.1a) that can be flexibly shaped to influence sensory processing in a goal-

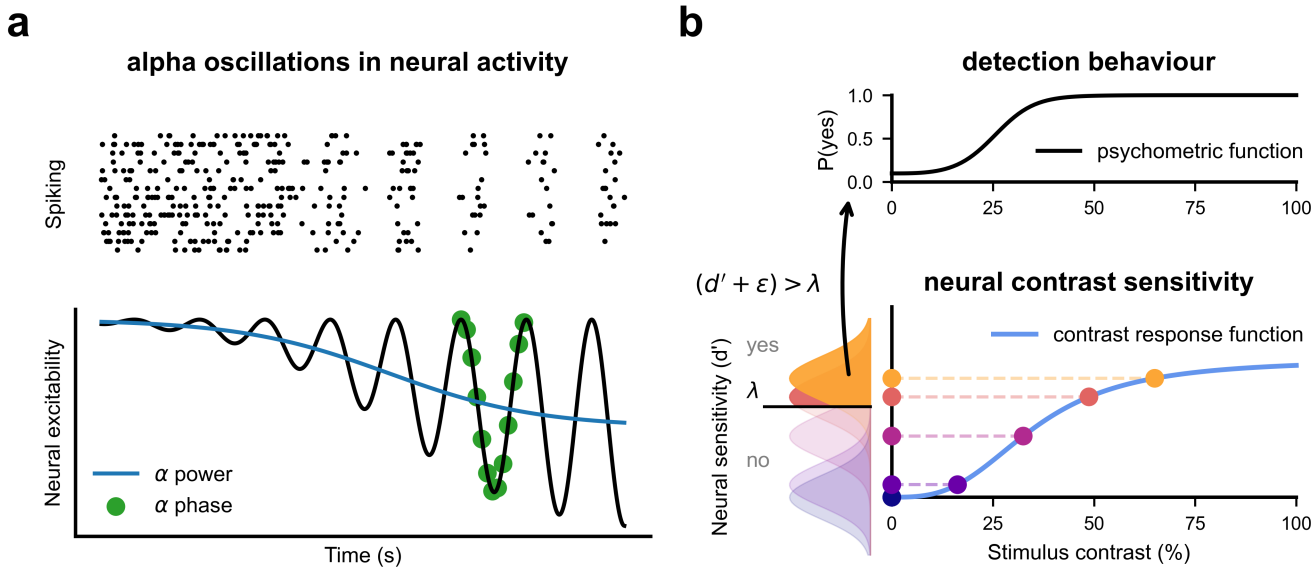

Figure 2.1: Theoretical relationship between neural activity, alpha oscillations, and perception

a Rhythmic suppression of neural spiking activity at ~10 Hz (upper panel: raster plot showing the firing of a population of simulated neurons). The plot shows rhythmic inhibition emerging gradually toward the right. Stronger pulses of inhibition spaced 100 ms apart result in synchronous firing. Lower panel shows inhibitory pulses (black line) in population level activity reflected in local field potentials or scalp EEG. Pulsed-inhibition underlies measures of alpha-band activity such as power, which increases as a consequence of stronger inhibitory pulses (blue line), and phase (green dots). **b** According to Signal Detection Theory, neural sensitivity to stimulus contrast intensity produces a psychometric function in detection (upper panel; x-axis: contrast; y-axis: probability of detection). While the psychometric function is observed experimentally, performance is determined by latent neural responses (lower panel). The coloured dots show evenly-spaced contrast values that produce neuronal responses of varying strength (due to a nonlinear transducer function in the early visual system; blue line). Neural noise shifts response strength (d') along the y-axis (lower panel) and produces a distribution across trials. A decision rule asserts that a detection response is made when the neural response (plus noise, ϵ) exceeds a criterion level (horizontal black line, λ). The psychometric function thus reflects the probability that noisy neural responses exceed this criterion as a function of contrast.

dependent manner (Klimesch, 2012; Klimesch et al., 2007; Mathewson et al., 2011; Mazaheri & Jensen, 2010; Schalk, 2015).

In humans, alpha oscillations are readily measured using non-invasive methods with high temporal resolution, such as electroencephalography (EEG) and magnetoencephalography (MEG). The phase of spontaneous alpha oscillations predicts perceptual reports of weak, near-threshold visual stimulation (Busch et al., 2009; Dugué et al., 2011; Harris et al., 2018). Phasic modulation of visual detection also depends on strong alpha amplitude (Alexander et al., 2020; Mathewson

et al., 2009), a relationship supported by causal evidence from Transcranial Magnetic Stimulation (TMS) over occipital cortex (Fakche et al., 2022). These findings are consistent with a pulsed functional inhibition hypothesis (Jensen & Mazaheri, 2010; Mathewson et al., 2011) in which the strength of phasic modulation depends on the amplitude of the spontaneous oscillation. That is, inhibitory brain states differ mostly strongly from peak to trough during high amplitude oscillations, and this phasic modulation is diminished during weak oscillations. However, the computational mechanisms by which pulsed inhibition rhythmically modulates perception remain unclear.

Signal Detection Theory (SDT; Tanner and Swets, 1954) provides a powerful tool with which to investigate the computations underlying fluctuations in visual sensitivity. In the context of visual neuroscience, SDT provides a generative model of visual behaviour based on latent neural responses and a decision rule (Fig. 2.1b). Early studies that linked alpha phase to visual perception did not use SDT methods (Busch et al., 2009; Mathewson et al., 2009) and thus could not separate perceptual from decisional influences of alpha phase on behaviour, as would be possible with SDT. Recent studies have used SDT to relate changes in alpha power (alpha phase was not measured) to fluctuations in decision bias, rather than sensory sensitivity (Iemi & Busch, 2018; Iemi et al., 2017; Limbach & Corballis, 2016). Critically, however, these studies only assessed performance at a single contrast level. Neural sensitivity is known to vary non-linearly with stimulus intensity (Albrecht & Hamilton, 1982), meaning that perception could exhibit different patterns of modulation depending on the stimulus intensity range examined. Testing the full range of stimulus intensity values is therefore required to gain a complete picture of how neural oscillations relate to perception.

Distinct computational mechanisms can shape the psychometric function in characteristically different ways, such as the response gain and contrast gain computations that have been proposed for attentional modulation of neural responses (e.g. Reynolds and Heeger, 2009). Only two studies have sought to relate alpha power to visual sensitivity using a full range of stimulus intensities and yielded separate support for response gain (Chaumon & Busch, 2014) or baseline/criterion modulation (Pilipenko & Samaha, 2024). Critically, both studies modelled psychometric functions using detection rates, instead of underlying SDT parameters, and neither examined the association between these mechanisms and the phase of alpha oscillations. A clear understanding of how alpha oscillations shape visual processing, via pulsed inhibition, requires the use of SDT analyses across the com-

plete stimulus intensity range, combined with an analysis of the joint influence of alpha power and phase.

In the present study, we investigated the combined interactive effects of alpha power and phase on visual judgements. We used EEG to measure prestimulus alpha power and phase, and combined these into a modelled inhibitory dynamic that is consistent with the pulsed-inhibition hypothesis of alpha oscillations (Jensen et al., 2014; Mathewson et al., 2011; Mazaheri & Jensen, 2010; Schalk, 2015). We tested observers in both visual detection and discrimination across a wide contrast-intensity range, and we developed a pulsed-inhibition model that links oscillatory phase and power to latent sensory responses using SDT. As described below, alpha-based inhibition modulated the response gain of visual sensitivity in both tasks, whereas a baseline shift in excitability was evident only in detection. These results support a common perceptual mechanism through which alpha power and phase, together, regulate visual sensitivity. Critically, we found that the effect of phase within the pulsed-inhibition model was weaker relative to the power-related suppression of sensory responses, a previously unseen power–phase dynamic that is not captured by pulsed-inhibition models of rhythmic perception.

2.1 Results

Prestimulus Alpha Oscillations Modulate Visual Responses

To investigate the effects of alpha oscillations on visual perception, we recorded EEG while participants made visual judgements in both detection and discrimination tasks (see Methods). Participants (n = 8 x 4 sessions, per task) viewed brief (8 ms) oriented gratings of variable contrast intensities. In the detection task, participants indicated the presence or absence of the target within a set of spatial markers in the visual periphery (see Fig. 2.2a). In the discrimination task, participants judged whether the stimulus was spatially offset either to the left or right of the markers. Stimulus visibility was manipulated using contrast intensity, which was adaptively set in each session to both probe performance around threshold level and best estimate the full psychometric function over contrast. Behavioural performance improved as a function of stimulus contrast in both tasks, shown by increasing rates of detection (Fig. 2.2b) and greater accuracy of location discrimination (Fig. 2.2c) with higher contrast. We used independent component analysis to isolate (per session) a source of posterior scalp EEG activity in the contralateral hemisphere, relative to the stimulus, that responded to stimu-

lus presentation. In both tasks, the magnitude of the stimulus-evoked responses was driven exogenously by stimulus contrast (Fig. 2.2d) supporting the source components' relevance for visual processing.

We next examined whether the evoked visual responses were modulated endogenously by spontaneous alpha-band activity. Prestimulus oscillations were detected on a single-trial basis using a spectral parameterisation algorithm to separate oscillations from background aperiodic activity (Donoghue et al., 2020), using data in a 500 ms window prior to stimulus onset. Estimates of instantaneous power and phase were then obtained for detected oscillations in the 8–14 Hz alpha frequency range. In both tasks, alpha power significantly predicted the amplitude of peaks in the evoked neural activity. Specifically, the negative peaks at ~200 ms post-stimulus were weaker when alpha power was stronger at stimulus onset (*detection*: β = -0.066, $t_{(19345)}$ = 2.62, P = .009; *discrimination*: β = -0.106, $t_{(19345)}$ = 2.72, P = .007), consistent with previous work showing suppression of early event-related potentials by alpha ongoing oscillations (Iemi et al., 2019). Furthermore, alpha phase at stimulus onset significantly predicted peak evoked amplitudes in both tasks (*detection*: d_z = 1.287, P < .001; *discrimination*: d_z = 1.745, P < .001). Fig. 2.2e shows the evoked waveforms plotted as a function of phase. In summary, the modulation of evoked responses by both power and phase demonstrates that neural responses to the stimuli were modulated endogenously by the state of spontaneous alpha activity at stimulus onset.

Next, we asked whether alpha power and phase impacted perception by modulating behavioural performance across the two tasks. We first used a binning approach to quantify behavioural accuracy and to visualise performance modulation across contrast, shown in fig. 2.3a–b, consistent with past approaches ((Chaumon & Busch, 2014); see Methods). Performance was equated across sessions by computing contrast as a percentage of thresholds and behavioural accuracy was then binned by alpha power and phase within each contrast level. To compute a summary measure of performance modulation that aggregates over contrast-specific effects, the slope of alpha-related modulation was averaged across contrast bins. We found stronger alpha power significantly decreased 'stimulus present' responses in the detection task (mean slope across contrast bins [95% highest density interval (HDI)] = -0.005 [-.01, -0.001], bootstrapped P value (P_{boot}) = .005). Similarly, stronger alpha power led to significantly fewer reports of the correct target location in the discrimination task (mean slope = -0.005 [-0.01, 4e-5], P_{boot} = .024). Given that prestimulus phase modulated the evoked

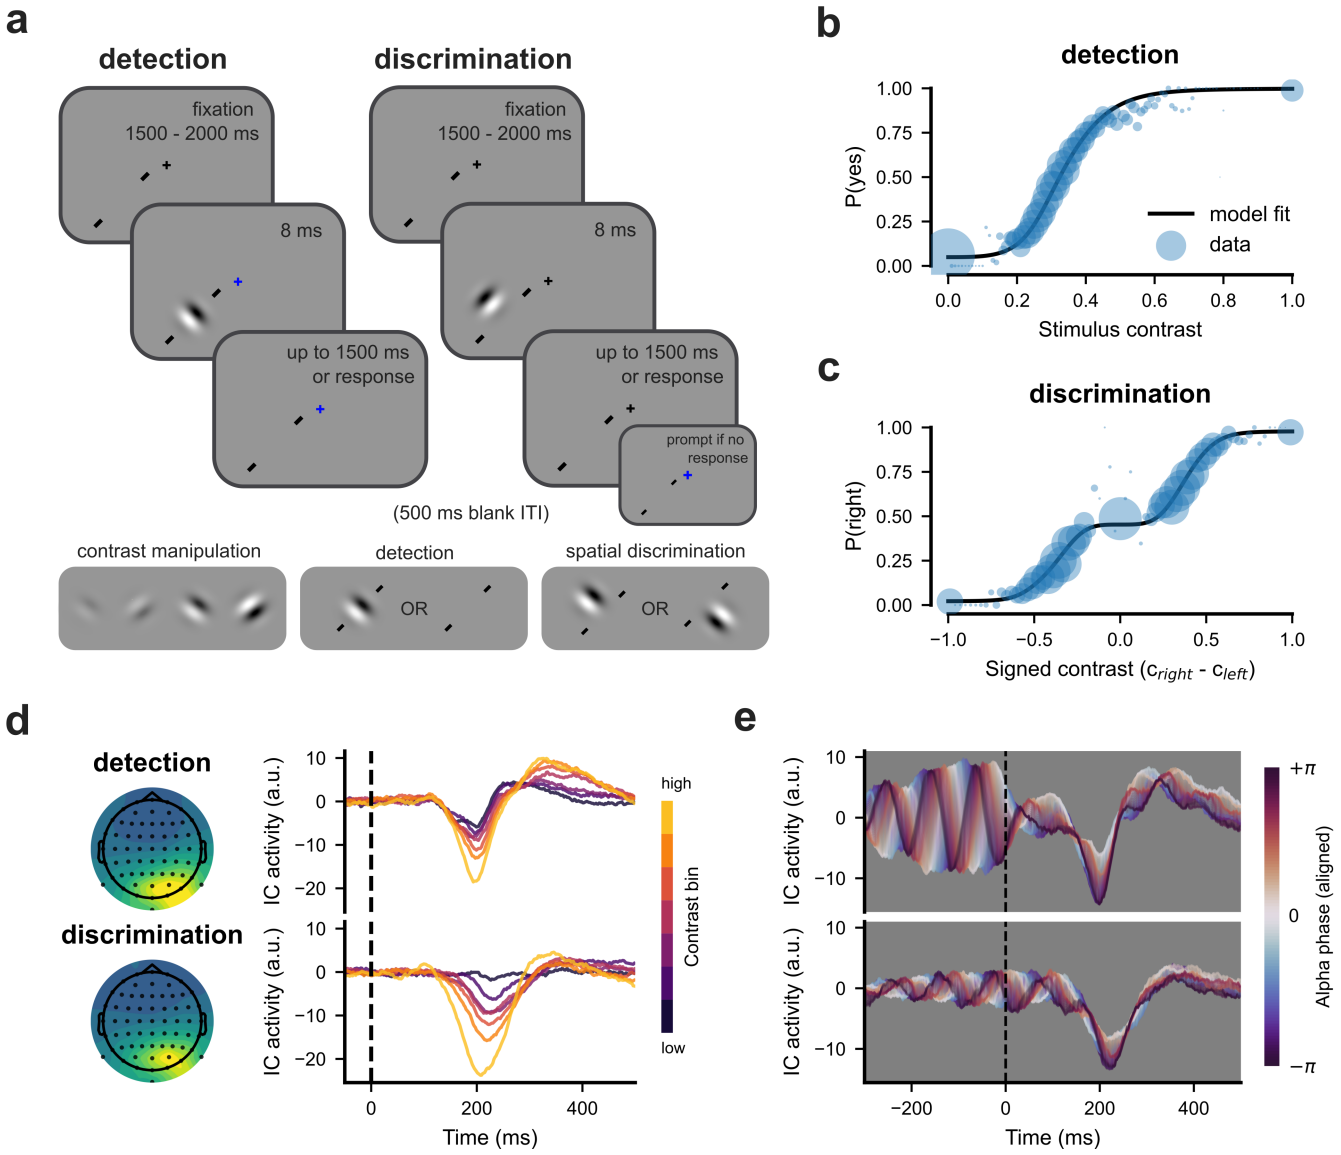

Figure 2.2: Behaviour and EEG responses in the two visual tasks

a Participants viewed grating stimuli that varied in contrast and, in separate experiments, reported stimulus presence/absence (detection; left-hand panels) or spatial position relative to onscreen diagonal tick markers (discrimination; right-hand panels). **b** Psychometric function of contrast-driven detection responses. Black line shows the predicted detection rate from a signal detection model, fitted hierarchically over participants. Circles show data averaged over participants, with sizes reflecting the number of trials presented at each contrast. **c** Psychometric function of spatial discrimination judgements over stimulus contrast (conventions as in **b**). The abscissa shows the signed contrast difference between the right (positive) and left (negative) spatial locations. The ordinate is the probability of reporting the stimulus in the rightward location. **d** Left panel: Scalp topographies of sensor activations from the analysed independent components (z-scored and averaged over sessions; stronger activation is shown brighter, arbitrary units). Right panel: Evoked EEG responses increased with stimulus contrast, brighter lines show increasing levels of binned contrast (dashed line shows stimulus onset time). **e** Evoked responses plotted by alpha-band phase at stimulus onset (estimated a sliding Von Mises window, $\sigma = 15°$, over phase values aligned, per session, to the phase with greatest evoked activation).

responses, phase values were circularly realigned relative to the phase that predicted the strongest evoked response. Behavioural responses were then binned by the circular distance from this optimal phase. However, phase did not influence overall responses, collapsed across contrast bins, in either task (*detection*: mean slope = -0.003 [-0.007, 0.002], P_{boot} = .147; *discrimination*: mean slope = -0.001 [-0.006, 0.005], P_{boot} = .425). Thus, overall behavioural accuracy was substantially modulated by alpha power, but not by alpha phase.

Behavioural Accuracy Modulation by Alpha Power and Phase

To investigate how alpha power and phase might modulate behaviour in a way that interacts with stimulus contrast, we conducted logistic regression analyses to predict behavioural accuracy as a function of contrast (see Methods). In the detection task, the addition of both power and phase predictors significantly contributed to the model likelihood relative to a null behaviour-only model (nested likelihood ratio tests; *power*: $\chi^2_{(4)}$ = 4,690.9, P < .001, AIC: 24,175, BIC: 24,248; *phase*: $\chi^2_{(4)}$ = 4,680.5, P < .001, AIC: 24,185, BIC: 24,258; *null*: AIC: 28,858, BIC: 28,899). This was also true for the discrimination task (*power*: $\chi^2_{(4)}$ = 2,654.1, P < .001, AIC: 20,081, BIC: 20,151; *phase*: $\chi^2_{(4)}$ = 2,652.5, P < .001, AIC: 20,082, BIC: 20,153; *null*: AIC: 22,727, BIC: 22,767). The interaction between power and phase added significantly to model likelihood in detection ($\chi^2_{(8)}$ = 20.789, P = .008, AIC: 24,180, BIC: 24,318) but not for discrimination ($\chi^2_{(8)}$ = 14.272, P = .075, AIC: 20,084, BIC: 20,218). In both tasks, the best negative log-likelihood was found for the power–phase interaction models (shown in Fig. 2.3c&d). However, information criteria that include a penalty for model complexity, such as AIC and BIC, consistently favoured the power model more strongly in both tasks (Supplementary Fig. 1). Together, these results suggest that both power and phase modulate visual performance, with an interaction effect emerging in the detection task, but there is stronger evidence for power as a model predictor of behavioural modulation.

Pulsed-Inhibition Alters Response Gain and Baseline Activity

Next, we investigated how alpha oscillations modulate the perceptual mechanisms that give rise to the specific shape of the psychometric function. We used a generative signal detection model to explain behaviour in terms of the latent sensory responses that underlie detection and discrimination choices. The strength of these

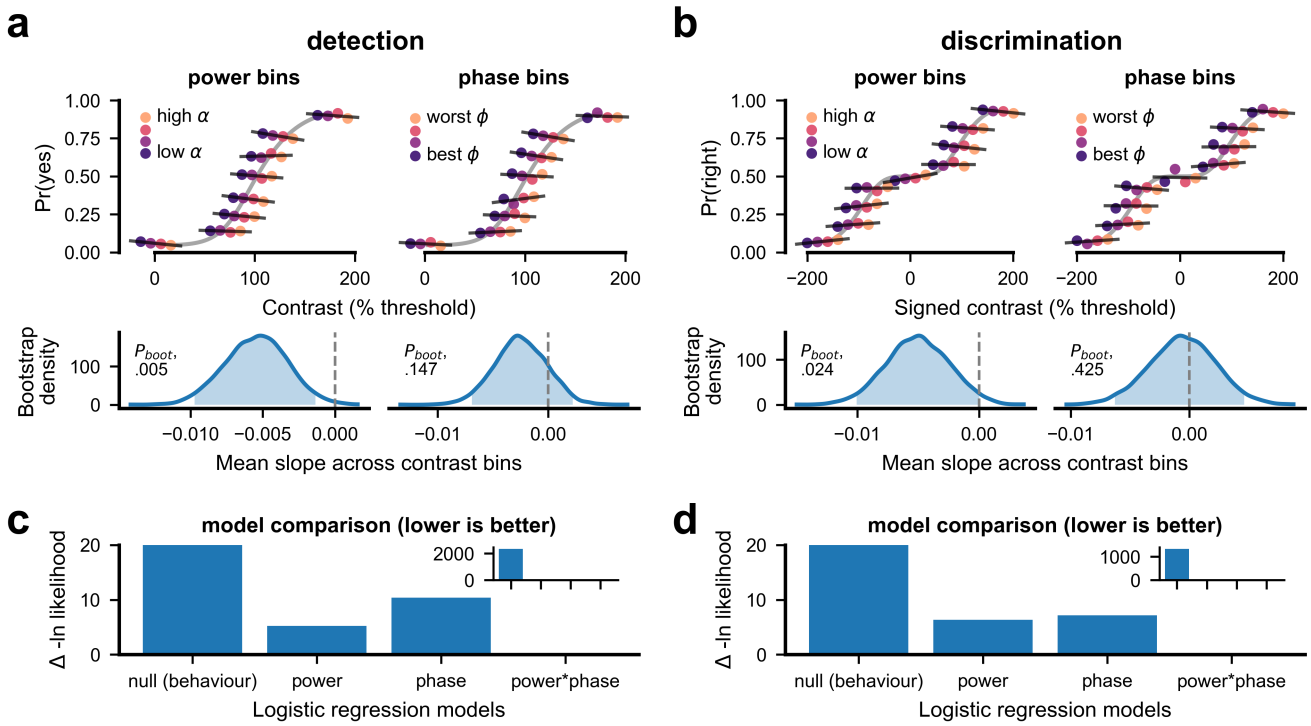

Figure 2.3: Behavioural accuracy by power and phase

a–b Binning analyses of accuracy for the detection (**a**) and discrimination tasks (**b**). Upper panels show accuracy binned by stimulus contrast for power (left) and phase (right). Darker colours show power/phase bins where theory suggests performance should be improved. Contrast jitter (i.e. spreading across the x-axis within each group of points) was added to facilitate visual inspection of power/phase differences within each contrast bin. The dark lines show the fitted slope of accuracy modulation within a contrast bin. Lower panels show the density of the average modulation slope across bins (estimated using kernel density estimation based on 5,000 bootstrap samples). Inset text shows the probability of samples lying below zero (dashed gray line). **c–d** Comparison of logistic regression models fitted to detection (**c**) and discrimination (**d**) data. Separate model fits were conducted for behaviour only (null), power, phase, and power by phase interaction terms. The y-axis (truncated) shows the change in negative log likelihood relative to the best fitting (power*phase) model, higher values indicate worse fit. Inset shows the full y-axis which captures the worse fit of the null model.

modelled sensory responses is a function of stimulus contrast (see Fig. 2.1b) and, we hypothesise, may also be impacted by alpha-related inhibition. Importantly, we combined alpha power and phase estimates into a model of pulsed-inhibition (in line with Fig. 2.1a; see *Pulsed-inhibition modelling* in the Methods) to compare the computational mechanisms by which inhibition shapes these latent sensory responses, and thus behaviour.

We fitted four models that describe the potential sensory modulation by pulsed-inhibition (see Fig. 2.4a–b, left columns). First, inhibition might affect

visual sensitivity via *contrast gain* which produces poorer sensitivity to weaker contrast values and is reflected in a rightward shift in the psychometric functions. Second, inhibition could divisively scale visual responses, via *response gain*, without changing the underlying range of contrast that the visual system is sensitive to. Third, *noise* in visual responses might be altered by inhibition, which affects performance mostly at the lowest and highest contrast levels. Finally, inhibition might affect the strength of visual responses regardless of contrast, via *criterion* modulation, which affects baseline performance in the psychometric function (i.e. zero-contrast performance in detection, or biased location choice in spatial discrimination).

In the detection task, we found the clearest support for pulsed-inhibition of sensory response gain (Fig. 2.4a; standardised parameter estimate (β_z [95% HDI]) = -3.388 [-5.413, -1.253], Bayes Factor (BF_{10}) = 20.27). Additional support was found for pulsed-inhibition of criterion (β_z = -2.814 [-5.497, -0.191], BF_{10} = 5.49), in line with previous findings (Pilipenko & Samaha, 2024). We found no support for oscillations affecting either contrast gain (β_z = -0.899 [-2.113, 0.227], BF_{10} = 0.73) or sensory noise (β_z = 0.167 [-0.79, 1.087], BF_{10} = 0.62). In the discrimination task, we found additional support for response gain modulation (β_z = -2.499 [-4.256, -0.48], BF_{10} = 3.16), but not for any of the remaining models (*contrast gain*: β_z = -1.509 [-3.382, 0.3], BF_{10} = 1.03; *noise*: β_z = 1.283 [-0.325, 3.07], BF_{10} = 0.89; *criterion*: β_z = 0.556 [-0.127, 1.339], BF_{10} = 0.66). Together, the modelling results suggest that, across tasks, alpha inhibits visual sensitivity to contrast via response gain modulation, and via baseline neural activity in the detection task but not in the discrimination task. We note the latter result is explained by the two-alternative forced-choice design of the discrimination task, which forces changes in perceptual bias to be matched between choice options and to thereby cancel out.

Perception is Modulated by Both Sustained and Pulsed Inhibition

We next examined the relative effects of alpha power and phase within our pulsed-inhibition model of perception. A key aspect of the pulsed-inhibition model is that the strength of phasic perceptual modulation depends on power (Fakche et al., 2022). To capture this in our signal detection modelling, alpha power and phase pair-values are mapped onto a single inhibitory value that is used to predict latent sensory modulation on each trial. The mapping to inhibition is shown in the upper panel of Fig. 2.5a and is increasing with power and pulsed across phase

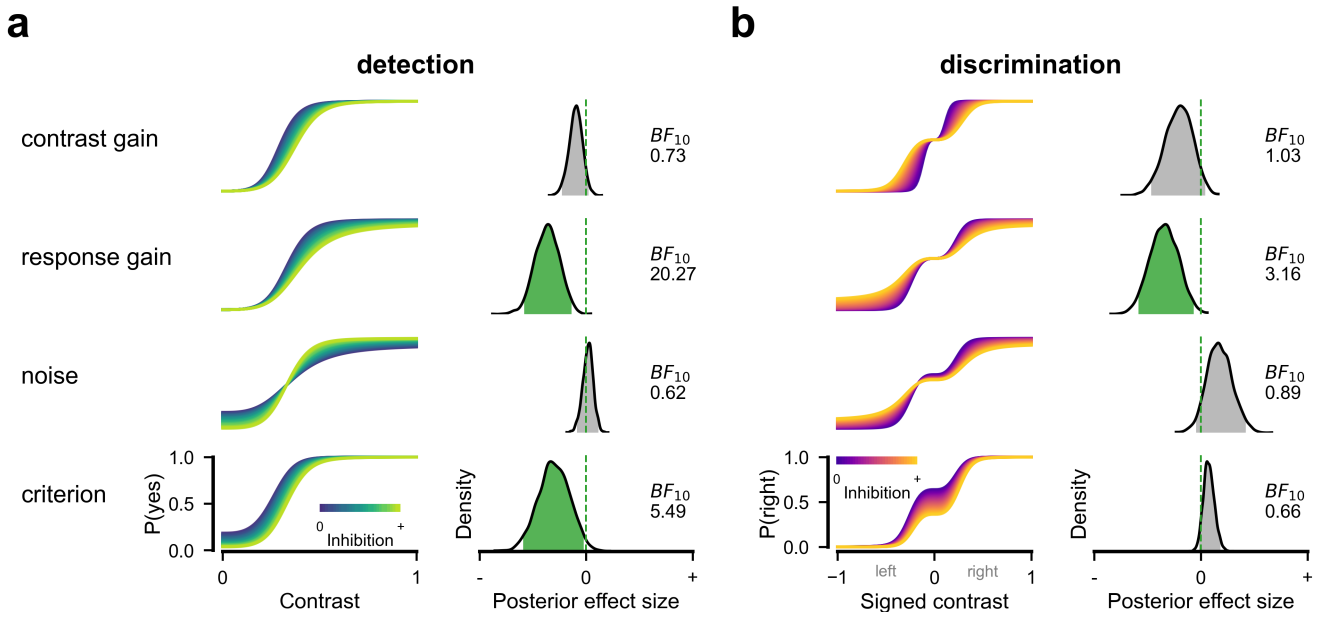

Figure 2.4: Model-based effects of alpha inhibition on the psychometric function

a Signal detection modelling in which pulsed-inhibition modulates the detection psychometric function via four putative mechanisms: contrast gain, response gain, (sensory) noise, and criterion. Left column: Simulated modulations of each parameter in the psychometric function under increasing inhibition (lighter colours). In detection, the y-axis is the probability of a 'yes' response. Right column: The inferred posterior densities of the modulation effects. To facilitate comparison, parameter values were standardised using the standard deviation among the hierarchically fitted sessions. Shaded regions of the distribution show the 95% Highest Density Interval and are shown in green when this does not include zero (shown by green vertical dashed line). Text insets show the Bayes Factor in support of the alternative. **b** Inhibition models for the discrimination task. Conventions as in **a** apart from the left column, where simulations of the psychometric function for discrimination are plotted over signed contrast (x-axis) and show the probability of reporting the stimulus as occupying the right-side location (y-axis). The underlying effects within the signal detection models do not change between tasks, but the psychometric function is nonetheless different because of the two-choice response structure.

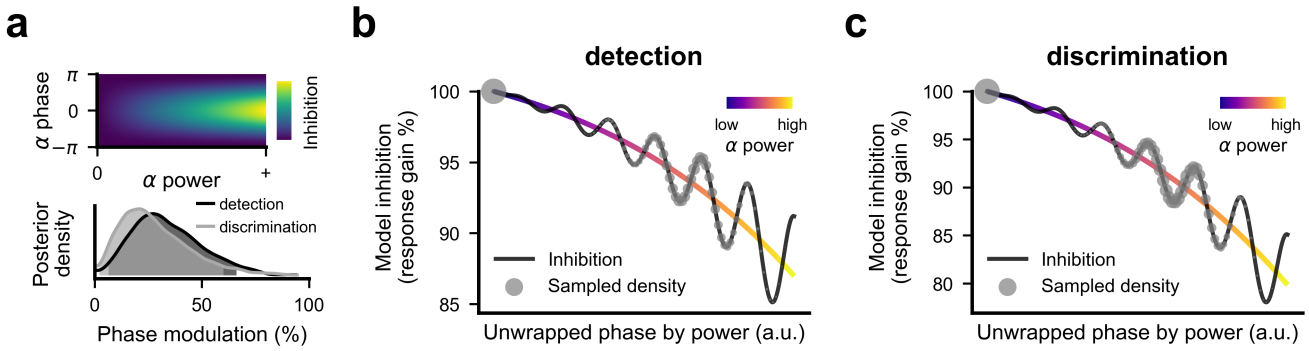

Figure 2.5: Modelled profile of behavioural inhibition via power and phase

a Upper panel: Alpha power and phase values are mapped to an inhibitory value, shown as a two-dimensional space. Lighter colour reflects stronger model inhibition that grows with higher alpha power (x-axis) and is dependent on phase (y-axis). Lower panel: The posterior density of a parameter controlling the percentage extent of phasic modulation. Zero values reflect no contribution of phase corresponding to power modulation only. **b–c** Reconstructed response gain modulation in the detection (**b**) and discrimination (**c**) tasks. The black line shows the model-predicted change in response gain (y-axis) plotted at a continuously unwrapped alpha phase and alpha power (coloured line). Note this relationship is idealised as not all alpha phase-power combinations are shown, but the inhibitory profile is conveyed, nonetheless.

values, consistent with the pulsed-inhibition hypothesis (Mathewson et al., 2011; Mazaheri & Jensen, 2010). To examine the dependence of phase on power, we included an additional weighting parameter in the behavioural modelling that describes how much phase contributes to the inhibitory values used to predict latent sensory modulation (fitted freely during the above SDT model inferences). Fig. 2.5a (lower panel) shows the posterior probabilities for this parameter and reveals that the contribution of phase to behavioural modulation was relatively weak in both the detection (mean weight [95% HDI] = 34% [6.4, 66.2]) and discrimination task (mean weight = 28.3% [2.1, 60.1]). We further visualised the inferred profile of the pulsed-inhibition of response gain in both tasks (see Fig. 2.5b–c). This shows that the fitted model supports a sustained suppressive effect at high power that is not fully released at less inhibited phases, as pulsed-inhibition theories suggest it should. Together, these analyses suggest that alpha power produces a sustained suppressive effect on behaviour, alongside weaker phasic modulation that does not return to unsuppressed behavioural performance across the phase cycle.

2.2 Discussion

We investigated how human visual sensitivity is shaped by spontaneous alpha-band neural oscillations. Using measures of power and phase at the onset of visual stimulation, we quantified their influence upon psychometric performance in both detection and discrimination tasks. By developing a novel modelling approach, we were able to link power and phase into a combined quantity that reflects an 'inhibitory drive' in line with predictions arising from pulsed-inhibition theories (Mathewson et al., 2011; Mazaheri & Jensen, 2010). We then inferred the trial-by-trial effects of inhibitory drive upon latent sensory responses within a signal detection theoretic paradigm. Our results show that this combined alpha power–phase predictor is associated with the suppression of sensory evidence, shaping behavioural performance across the full range of presented contrast intensities. Moreover, we uncover the computational basis for oscillatory perceptual effects on sensitivity, finding a response gain mechanism that covaried with alpha oscillations in both detection and discrimination tasks. We further demonstrate that the contribution of phase to the modulation of perceptual performance is far weaker than would be expected from a strong version of the pulsed-inhibition model. Specifically, when alpha power is high, sensory responses are reduced and rhythmically modulated by alpha phase, but there is no phasic return to the unsuppressed levels of behaviour seen when alpha power is low. These results suggest that alpha-frequency neural inhibition manifests in both sustained and weakly phasic suppression of response gain that, together, shape perceptual ability.

A critical feature of our modelling is that it accounts for the variable influence of phase that would be expected at different levels of oscillatory power. That is, the same phase values can be associated with different levels of inhibition depending on power, such that a trough at low power may be associated with perceptual effects that are equal to those of a peak at higher power (see Fig. 2.5). This dependence of phase effects on power could result in undetectable phase-modulation if not deconfounded. Such a fact has been ignored in most previous investigations of oscillatory phase-effects and may be a reason why some past work (e.g. the contrast binning analyses of Chaumon and Busch, 2014; recapitulated here, with results shown in Fig. 2.3) has apparently supported power- but not phase-modulation of perception. Studies that have acknowledged this issue have typically split their data into binary, high versus low alpha power conditions (e.g. Fakche et al., 2022; Mathewson et al., 2009; Ruzzoli et al., 2019), expecting to observe phase effects only in the high-power condition. Our analysis is robust to

these problems, as we model a continuous interactive power–phase relationship and can include all trials during model estimation. This approach can disentangle the possibly opposing effects of phase at different levels of power with no inherent loss of statistical power. More importantly, our combined model of alpha-based inhibition is coherent with theoretical proposals concerning the role of alpha oscillations in neural sensory processing (Klimesch et al., 2007; Mathewson et al., 2011; Mazaheri & Jensen, 2010). These theories suggest that the influence of phase depends intrinsically on oscillatory power, due to their common relation via a pulsed inhibitory effect upon neural activity.

It is well established that alpha oscillations inhibit neural spiking (Dougherty et al., 2017; Haegens et al., 2011), but the computational and perceptual consequences of this suppression are currently unclear. Much prior work investigating links between alpha and perception has been unable to speak to the mechanisms involved as it focussed exclusively on hit rates in detection tasks (e.g. Busch et al., 2009; Harris et al., 2018; Thut et al., 2006). One challenge has been the need to disentangle the influence of alpha on perceptual sensitivity from changes in bias (which may have either perceptual or decisional sources; for review, see Samaha et al., 2020). This requires probing behavioural modulation using target-absent trials to allow estimation of the false-alarm rate. Studies employing this approach have typically found that alpha power modulates false alarm rates to the same extent as it modulates successful threshold-level target detections, suggesting a change in the criterion for detection (Iemi et al., 2017, 2019; Limbach & Corballis, 2016). This approach presents another challenge, however, as performance may be modulated differentially outside the contrast range used to probe threshold-level performance. Studies that have modelled performance over the full contrast range have yielded conflicting results. Chaumon and Busch (2014) concluded that alpha power modulation was associated with changes in response gain. In contrast, Pilipenko and Samaha (2024) observed a vertical shift in the contrast response function with alpha power that they interpreted as evidence of a change in criterion. However, neither of these studies employed SDT in their modelling of the complete contrast psychometric function and instead infer changes in the psychometric function using a hit-rate axis, which can misrepresent nonlinear effects on sensitivity. By estimating the source of alpha modulation across varying levels of contrast in a latent signal-detection model, we were able to show that perceptual sensitivity is affected in a manner consistent with changes in response gain across both discrimination and detection tasks. This suggests that the inhibition

associated with alpha oscillations divisively scales neural responses depending on the strength of sensory input and thus alters perceptual sensitivity.

In addition to response gain modulation in both tasks, we observed a robust alpha-related change in criterion that was unique to the detection task. This result is consistent with past findings (Iemi et al., 2017, 2019; Limbach & Corballis, 2016) and suggests that alpha-frequency inhibition modulates baseline firing rates in visual cortex (Samaha et al., 2020). This modulation of baseline firing alters the probability of reporting a stimulus as present on the basis of noise excitability in the visual system (i.e. false alarms). It will, however, have little influence on discrimination performance, if perceptual noise is equally likely to affect both response options, and so result in no overall change in bias (Iemi et al., 2017). In line with this, we did not observe a criterion effect in our discrimination task. Together, the effects of response gain and criterion suggest a mechanism by which alpha-based inhibition scales visual neural activity without altering underlying sensitivity to contrast (i.e. via a threshold or contrast gain change). This would be consistent with inhibition of visual cortex that scales all components of cortical firing activity, including both baseline neural activity and the afferent stimulus-driven responses from earlier neurons.

Our findings suggest that over the cycle of the alpha oscillation, alpha phase is associated with similar mechanistic fluctuations as alpha amplitude. Critically, this was true across both tasks. These results provide a mechanistic explanation for previous findings showing fluctuations of detection across the alpha cycle (Busch & VanRullen, 2010; Busch et al., 2009; Harris et al., 2018; Mathewson et al., 2009; Zhou et al., 2021), and align with recent results that indirectly suggest alpha phase and amplitude are associated with the perceptual consequences of lateral inhibition in early visual regions (Williams et al., 2024). Importantly, we also observed direct evidence for suppression of visual responses as alpha power and phase both reduced the magnitude of stimulus-evoked neural response (see also, Busch and VanRullen, 2010; Dou et al., 2022; Iemi et al., 2019; Mathewson et al., 2009). These results provide converging evidence for the suppression of early visual responses by alpha-frequency neural oscillations (c.f., Zhigalov and Jensen, 2020).

Our pulsed-inhibition model is a direct implementation of the theorised relationship between alpha-frequency oscillations and the underlying neural activity that is relevant for perceptual processing (Klimesch, 2012; Klimesch et al., 2007; Mathewson et al., 2011; Mazaheri & Jensen, 2010). Signal detection modelling also

allowed us to compare the relative contributions of alpha power and phase toward the modulation of perception. Theory suggests these should be equivalent; the effect of power should simply be the effects of phase, at that power level, integrated over the course of an oscillatory cycle. However, we observed that the relative effect of phase in this model was only ~30% of what theories suggest it should be, across both tasks. Such a finding may provide a plausible explanation for why phase effects have been difficult to observe in some scenarios (e.g. Ruzzoli et al., 2019; see also Harris, 2023). The source of this discrepancy is difficult to pinpoint. Electrophysiological features that were not included in our modelling, such as the non-sinusoidality of posterior alpha (Cole & Voytek, 2017) or saturation of modulation at much higher power levels, could weaken an observed phase-power correspondence. However, our results suggest the weakened influence of phase was present even at moderate levels of power, so it would be surprising for either of these possibilities to reduce the observable phase effect by ~70%. Instead, our results suggest the existence of a tonic, non-pulsatile suppressive effect that scales with alpha amplitude. One candidate mechanism are GABAergic inhibitory networks that are known to be associated with alpha oscillations (e.g. (Lozano-Soldevilla et al., 2014); Baumgarten et al., 2018) and commonly have time-constants longer than the ~100 ms cycle of an alpha oscillation. The activation of GABA-B receptors in visual cortex (and elsewhere) can produce inhibitory effects spanning hundreds of milliseconds (Wilson et al., 2025), suggesting a possible mechanism by which alpha oscillations could produce tonic inhibition of visual processing.

In summary, here we used a model-based analysis of visual behaviour to reveal the mechanistic bases by which measured neural oscillations in the alpha band shape perceptual sensitivity. Our results suggest that sensory evidence is divisively suppressed, in line with the idea that alpha oscillations reflect an inhibitory influence upon neural processing. We show that the strength of phasic modulation is not fully consistent with a strong version of the pulsed-inhibition model. Rather, our results suggest the effects of alpha phase may be weaker than the largely tonic effects of alpha power on perceptual behaviour. These findings underscore the necessity of modelling brain activity and perception jointly to explain the nonlinear interactions that link alpha-band dynamics to perceptual behaviour.

2.3 Methods

Participants

Nine healthy adults with normal or corrected-to-normal vision participated in the study (mean age, 26.9 years; range, 22–40 years; five females), including two non-naive authors HB and AH. Experiment 1 (detection task) involved eight participants (4 males and 4 females). Experiment 2 (discrimination task) involved the same participants, but one female dropped out and was replaced by another female participant. No participants were excluded from any analyses. The experiment was approved by The University of Queensland Human Research Ethics Committee. All participants gave written informed consent and all but the authors were compensated at a rate of $20/h AUD.

Apparatus

The stimuli were presented on a gamma-corrected 24-inch LCD monitor (VIEWPixx 3D, VPixx Technologies) with 1920 x 1080 display resolution and 120 Hz refresh rate within a dark, electromagnetically shielded room. Viewing distance was maintained at 58 cm using a chin rest. Stimuli were generated using MATLAB (MathWorks, R2020a) with the Psychophysics Toolbox v.3.0.17 (Brainard, 1997; Kleiner et al., 2007). A Biosemi ActiveTwo system (Biosemi, Amsterdam, Netherlands) was used to record 64 Ag-AgCl electrodes digitised at 1024 Hz and arranged in the standard 10-10 layout (Oostenveld & Praamstra, 2001). Per Biosemi design, the Common Mode Sense and Driven Right Leg electrodes served as reference and ground. Eye muscle activity was monitored using two EOG electrodes placed above and below the right eye, and two placed at the outer canthi of each eye.

Stimuli, Task, and Experimental Procedure

The stimuli were sinewave gratings (4 cpd; ±45° random orientation; randomised phase) presented within a Gaussian window (0.5 dva sigma) on a mid-grey background. A centrally positioned black cross (size: 0.2×0.2 dva; line width: 2 pixels) was presented throughout the trial to reduce eye movements. The target stimulus was presented at a single lower left location on all trials (7.07 dva eccentricity, 5 dva below and 5 dva to the left of fixation) to maximise trial numbers for analyses without needing to counterbalance analyses across hemispheres. In the detection task (experiment 1) the stimulus was presented within two spatial markers (line

lengths: 0.2 dva; line width: 2 pixels) positioned along the 45° diagonal and offset from the centre of the stimulus by 2 dva (4 × sigma). The observers' task was to report whether the target stimulus was present or absent on each trial. In the spatial discrimination task (experiment 2), the stimulus was presented 1 dva along the opposite diagonal perpendicular to the spatial markers, and the observers task was to report whether the stimulus appeared to the left or right of the spatial markers (Fig. 2.2a).

Each trial started with a variable fixation period of 1500–2000 ms. The stimulus was then presented for 8.33 ms and observers had 1500 ms to make a response. In the detection experiment, the colour of the fixation cross was changed from black to blue at the stimulus onset, to remove temporal uncertainty. In the discrimination experiment, the fixation colour change occurred 1500 ms after the stimulus, to inform the participant that the stimulus was no longer displayed and a two-alternative forced choice response was required. Participants completed an hour-long introductory behavioural testing session to gain familiarity with each task prior to EEG recordings. Each participant then completed four EEG sessions of approximately 400 trials, resulting in 29,090 trials for analysis in experiment 1 and 22,020 trials in experiment 2. Breaks were given every few minutes.

Psychometric functions were fit online using a Bayesian approach to continually update a parameter posterior using a discrete grid approximation. We used a cumulative Gaussian psychometric function with parameters for the slope and location in log contrast, as well as independent upper and lower asymptotes. On the first 150 trials in each session, and a random 10% of trials thereafter, the stimulus contrast was adaptively determined using a look-ahead algorithm to find the contrast value that minimised the expected entropy of the parameter posterior (Watson, 2017). This allowed us to present contrast values that were most informative for quick estimation of the psychometric function parameters, and to ensure that this estimate continued to be correct throughout the session, as needed. On the remaining (majority) of trials, contrast was sampled from a normal distribution in log contrast centred on the current best estimate of the participants' threshold and with standard deviation of the inverse slope estimate. This ensured that the majority of stimulus values were sampled around the most informative parts of the psychometric function, where performance changes rapidly with contrast (Kontsevich & Tyler, 1999). In the detection experiment, an additional random 10% of the non-adaptive trials were chosen for zero-contrast stimuli to enforce extra sampling of the false alarm rate for signal detection analyses.

Electroencephalography

Recording and Preprocessing

Offline processing of EEG data was performed using EEGLAB v2024.2 (Delorme & Makeig, 2004). Signals were downsampled to 512 Hz, high-pass filtered at 0.5 Hz using *pop_eegfiltnew.m*, and re-referenced to the average of electrodes. Bad channels were automatically identified using the adaptive channel scoring algorithms from the FASTER pipeline (Nolan et al., 2010) and were interpolated using spherical splines. Independent component analysis (ICA) was then performed with the extended Infomax algorithm (Lee et al., 1999) using *pop_runica.m* on a copy of the data that was 1 Hz high-pass filtered, demeaned across epochs (Groppe et al., 2009), and automatically cleaned of noisy epochs using the *pop_autoreg.m* function with default parameters. Fitted ICA weights were then used to project the original data and epochs into the component space.

We used a semi-automated procedure to select focal components for further analyses. First, all components were ranked according to the unsigned correlation between the component activation maps and a posterior region of interest contralateral to the stimulus (electrodes Oz, O2, POz, PO4, PO8, P2, & P4). We then selected a single component from the top five ranking components. This was done by manually inspecting the trial-averaged prestimulus spectral density and the poststimulus evoked waveform for each component (see Supplementary Figures 2 and 3 for the component topographies, evoked responses, and frequency spectra of chosen components). The highest ranking component (commonly the first or second) that showed both an alpha spectral peak and a typical visually-evoked response was selected for analysis. To enforce sign consistency across participants and sessions, the sign of the component was flipped so that the maximally activated electrode was positive (this also ensured a common oscillatory phase interpretation across participants).

Oscillations

Prestimulus oscillatory power was estimated using a Fourier transform on the hamming windowed data within a -0.5–0 s prestimulus epoch (2048 sample padding). The *specparam* toolbox (Donoghue et al., 2020) was then used to separate aperiodic and non-aperiodic components within the power spectra. Frequencies between 1–40 Hz were used with the algorithm's default parameters, maximum peaks set to 4, and peak width limits set between 0.5–6 Hz (trial-

averaged R^2 = 0.81). From this, we extracted the peak frequency of detected oscillations within the 8–14 Hz frequency range at the single trial level. To then characterise alpha power and phase, we used the endpoint-corrected Hilbert transform method (Schreglmann et al., 2021) to estimate the instantaneous power and phase at stimulus onset. We applied this only on trials in which *specparam* detected an oscillation. This was done by band-pass filtering prestimulus epochs (-1.5–0 s) between +/- 1 Hz of the peak frequency, and then using the absolute value and angle of the resulting complex analytic signal as power and phase estimates, respectively.

Analyses of Evoked Responses

To characterise the magnitude of evoked responses, we detected the timepoint at which the (negative) peak minima was found in the poststimulus data. This was done separately per session using the averaged waveform of all trials in that session, and searching within a 160–240 ms window. To yield a single-trial measure of evoked amplitude, trial-level data were averaged within a 30 ms time window centered on the per-session peak timepoint. The time window used for automatic selection was determined by inspecting the waveform grand-averaged across all data.

A mixed effects General Linear Model (GLM) was used to predict the evoked response magnitude using alpha power. Random effects were fit for the intercept and slope terms. The effect of alpha phase was analysed using separate linear regressions, per session, of the sin and cos of phase. Session-specific analyses allowed the optimal phase of the regressed sinusoidal modulation effect to vary between participants and sessions, accounting for potential phase differences that stem from individual cortical differences and electrode placements in each session. We assessed group level significance by combining P values using Fisher's method. An effect size, d_z, was calculated for group-level phasic modulation by reconstructing the amplitude of phasic modulation from regression weights, and computing a z-score across sessions.

Behavioural Analyses

Data Transformation and Normalisation

Contrast, power, and phase data were normalised across sessions, allowing data to be collapsed together, while mitigating major subject-specific differences. Con-

trast was computed as a proportion of the threshold contrast that was adaptively located during each session. Alpha power was log transformed and mean-scaled, whereby non-zero power values were divided by the average power, per session, to express power in relative units (mean = 1). Given that we found significant group-level phasic modulation of the evoked response, we took the specific phase of these effects as the relevant phase for potential behavioural modulation. Thus, alpha phase was normalised by circularly subtracting the phase value that predicted the strongest evoked response magnitude, specific to each session. For analyses that model a linear phase effect (specifically, binning and logistic regression) phase was transformed by taking the absolute value, which yields an unsigned distance from the optimal evoked phase. Trials containing no detected alpha oscillation were assigned a normalised power value of zero and were also excluded from analyses that modelled phase exclusively (as phase is not meaningful in zero-power oscillations).

Binning Analyses

We used eight equally sized contrast bins, combining data across participants. Within each contrast bin the relevant alpha data (either power or phase) were further divided into four equally sized bins. Behavioural responses were then averaged within each contrast-power/phase bin. A first degree polynomial was then fit across the alpha power/phase bins to yield the slope of accuracy modulation within each contrast bin. We then took the average slope over contrast bins as an aggregate measure of behavioural modulation across contrast bins. This procedure was repeated 5,000 times using bootstrapped datasets. We calculated bootstrapped P values using the one-sided probability that samples lay above zero. That is, under the null hypothesis that alpha power/phase suppress behavioural performance, we would expect the bootstrapped aggregate slope samples to lie at or above zero.

In the discrimination task, behavioural choices were modelled as a function of signed contrast, using the difference in contrast between the right and left stimulus positions (Fig. 2.3b). This results in better accuracy when the probability of reporting the righthand position is higher for positive signed contrast and lower for negative signed contrast. We accounted for this by inverting the slopes of negatively signed contrast bins prior to aggregating across contrast bins. We included zero-contrast trials in an additional bin that was excluded when calculating the aggregate slope, as performance here is at chance level by definition (i.e. guessing).

Logistic Regression

Generalised Linear Mixed Effects models were used to assess the influence of alpha power and phase on behavioural accuracy, using MATLAB's *fitglme* with a logistic link function. Model fit was improved by including predictors for additional transformations of contrast, by raising contrast to exponents of 0.5 and 2. Including these terms models the nonlinear relationship between contrast and predicted behavioural accuracy, which is achieved by approximating curved changes on the linear predictor scale instead of utilising only the sigmoidal shape of the logistic link function (Knoblauch & Maloney, 2012). We included random intercepts in the model of the contrast terms and use this as a null model that fitted performance on the basis of stimulus contrast alone. Additional models were fitted using alpha power and phase, including intercepts for these terms as well as slopes that interacted with the contrast terms and thus describe how power and phase shape a contrast-related change in accuracy. Finally, we fit a power by phase interaction model. The significance of predictors was assessed using nested likelihood ratio tests to assess the contribution of increasing model complexity. This allowed us to infer the significance of power, phase, and power by phase interaction terms, by looking across the additional terms that interact with contrast. We also evaluated model performance by transforming information criteria (AIC, BIC) into weighted model probabilities (see Supplementary Figure 1; (Wagenmakers & Farrell, 2004)).

Signal Detection Modelling

We used a signal detection theory approach to model behavioural responses in both tasks. The stimulus was assumed to be represented in a one-dimensional sensory space, and its representational intensity varied on a trial-by-trial basis. This variation was modelled as a unit variance Gaussian distribution where the mean stimulus representation intensity is a function of the external stimulus contrast:

$$d'(c) = \frac{\kappa c^{\gamma}}{c^{\gamma} + \tau^{\gamma}} + \epsilon, \quad \epsilon \sim \text{Normal}(0, \sigma). \tag{2.1}$$

This function describes a saturating non-linear mapping of the stimulus contrast, c, onto an internal response strength (i.e. a *transducer function*) that has been studied in neurophysiology (Albrecht & Hamilton, 1982; Naka & Rushton, 1966; Sclar et al., 1990). The parameter κ controls the maximum strength of internal representation that reaches saturation, τ locates the contrast at which the function

becomes compressive and begins to saturate, and γ controls the rate at which stimulus contrast increases the internal response.

In the detection task, it was assumed that the observer responds with a yes decision on trials where this internal representation exceeds a fixed criterion, λ. A decision rule for the localisation task is explained further below. The probability of a binary yes-no decision is thus given by integrating the probability that the internal representation exceeds λ for a given stimulus contrast:

$$\psi_{yes}(c) = Pr(d'(c) > \lambda) = 1 - \Phi(\lambda - d'(c)) \tag{2.2}$$

where $\Phi(\cdot)$ is the cumulative Gaussian function, which performs an identical role to the probit link function within the General Linear Model (DeCarlo, 1998). The observed detection responses are modelled as a Bernoulli random variable where the probability of success is a given by the transducer function of contrast:

$$Pr(choice|c) \sim \text{Bernoulli}(\psi_{yes}(c)). \tag{2.3}$$

In the discrimination task, the observer chooses the spatial location of a single target stimulus (left or right of the spatial markers). We modelled this task as a 2-alternative forced choice (2AFC), in which the observer monitors both spatial positions and reports the location with the higher internal response on each trial. The probability of reporting the right-hand location is:

$$\psi_{right}(\mathbf{c}) = 1 - \Phi\left(\frac{\lambda - d'(c_{left}) - d'(c_{right})}{\sqrt{2}}\right). \tag{2.4}$$

In this context, the criterion λ is a bias parameter that shifts the decision in favour of a certain spatial interval but does not reflect a subjective threshold for the stimulus representations that are compared objectively. Note that we modelled the choice of a single response option, which ranged from 0–100% correct rather than 50–100% correct used for accuracy (Lesmes et al., 2015). We also modelled an additional lapse rate parameter in the discrimination task, which improved model fit and accounts for the larger number of guessing responses caused by additional temporal uncertainty in the task design. With lapse proportion, ϵ, the corrected response probability is $\psi'_{right}(\mathbf{c}) = \epsilon + [1 - 2\epsilon]\psi_{right}(\mathbf{c})$.

Pulsed-Inhibition Modelling

To combine alpha power, P, and phase, ϕ, values into a single scalar inhibitory value, I, we used the function on each trial, i:

$$I_i = P_i(1 + m\cos(\phi_i)). \tag{2.5}$$

This mapped single-trial power and phase into a value that was used as a covariate within the signal detection modelling. The weighting parameter $m \in [0, 1]$ controlled the extent of phasic modulation, with zero values providing no phasic modulation (i.e. a power-only modulation effect) and a value of one providing full phasic modulation (consistent with the functional form visualised in Fig. 2.1a). We tested four candidate mechanisms by which inhibition could modulate sensory processing. For response gain modulation, the output of the transducer function was scaled, using a log link function for inhibition modulation:

$$e^{gI_i} \times d'(c_i). \tag{2.6}$$

Here, the free parameter g controls the predicted modulation given inhibition and contrast values. Note that multiplying the output of the transducer function is mathematically equal to adjusting the κ parameter in 2.1. For contrast gain modulation, the transducer function threshold parameter, τ, was adjusted using a logit link for inhibition:

$$\tau_{adj} = \text{logit}^{-1}\left(\text{logit}(\tau) - CI_i\right) \tag{2.7}$$

where C is free parameter controlling contrast gain modulation. For noise modulation, the probability of task responses was modulated by divisively scaling the input to the cumulative Gaussian, G, in both ψ_{yes} and ψ_{right} defining

$$G_N(x) = G\left(\frac{x}{e^{NI_i}}\right). \tag{2.8}$$

Here, N is the free parameter controlling noise modulation via a log link for inhibition. Criterion modulation was achieved with the free parameter, b, by adjusting λ:

$$\lambda_{adj} = \lambda - bI_i. \tag{2.9}$$

We fit the Bernoulli response models described above to the single-trial data within a Bayesian framework using Hamiltonian Monte Carlo in PyMC v4 (Abril-Pla et al., 2023). Parameters were modelled using a mixed-effects approach to hierarchically model the variation across the 32 experimental sessions for each task (individual psychometric function fits are shown in Supplementary Figure 4). Here, participant-level contrast thresholds, τ, were fitted, so the untransformed contrast values were modelled. As d' is difficult to estimate at high values, we followed Lesmes et al. (2015) and fixed the parameter κ at 5. Before fitting the observed data, we conducted a prior predictive analysis to ensure that the choice of

priors was able to generate model predictions that were psychophysically plausible and consistent *a priori* with realistic data (i.e. monotonically increasing psychometric functions, without extreme/flat slopes, and without contrast thresholds near full/zero contrast). For the m parameter controlling phasic modulation, we used a flat Beta(1,1) prior. For the remaining modulation priors (b, g, N, C), we used weakly regularising Normal(0,2) priors.

Statistical Inference

Models were fit using the No U-Turn Sampling HMC method in PyMC. $\hat{R}$ values were less than 1.01 for all parameters, suggesting that sampling chains converged to the target posterior. The reported β_z values in text and figures refer to the standardised modulation parameters. These were obtained by dividing the posterior samples of the group-level modulation parameters (b, g, N, C) by their fitted standard deviations. Models were fit independently to facilitate interpretation, as parameters within the probit link are highly dependent on the specific values of all other parameters. For the detection modelling we verified that support for the response gain and criterion results remained when fit together, and we report this combined model in text and figures.

We conducted hypothesis tests using the inferred posterior probabilities of the standardised effects, and we assessed whether the 95% Highest Density Interval (HDI) of the posterior contained the null value (zero). We also performed Bayes Factor analyses to assess the strength of evidence for modulation. The fitted hierarchical models gave us insight into the variability of modulation effects (which we modelled as normally distributed across sessions), and this also provides information on the parameter value at the group level. We computed Bayes Factors using the posterior mean estimates of the standardised modulation effects. This was calculated on the basis that the standardised effect is expected to be zero under the null hypothesis (i.e. there is no effect relative to residual variability in the parameter) and follows a Cauchy distribution (with standard scale parameter of $\sqrt{2}/2$, Morey and Rouder, 2011) under the alternative hypothesis.

Data Availability

The behavioural and neural data generated in this study are available on the Open Science Framework database: https://osf.io/3rm8z/.

Code Availability

Code has been deposited on GitHub: https://github.com/henrybeale/alpha-power-phase-visual-sensitivity.

Chapter 3

Alpha-band Oscillations Modulate Sensory Noise in Efficient Visual Codes

A major task of the human visual system is to encode features of the external environment using limited neural resources. The efficient coding principle proposes that sensory systems are adapted to the statistical structure of the environment (Barlow, 1961). This framework has successfully explained key attributes of sensory processing, such as the non-uniform tuning of sensory neurons that is matched to environmental image statistics (Ganguli & Simoncelli, 2016) and distortions in perception that result from efficient internal representations (Wei & Stocker, 2017). In biological systems, however, neural activity is constrained by the brain's ever-evolving internal states. Neural oscillations are a prominent source of these endogenous dynamics, shaping moment-by-moment changes in cortical excitability and sensory responsiveness (Womelsdorf et al., 2014). Here we ask how visual coding is jointly shaped by the statistical structure of the external world and by the brain's own internal temporal states.

In the efficient coding framework, perceptual biases arise from internal representations that over-represent environmentally prevalent stimuli (Hahn & Wei, 2024; Wei & Stocker, 2015). An example are the cardinal vertical and horizontal orientations that dominate natural images (Girshick et al., 2011), which have more visual neurons tuned to them compared with oblique orientations (Li et al., 2003). This typically results in a bias in the perceived stimulus relative to the cardinals, such that perceptual estimates are viewed as further away from the

cardinals, or more tilted, than the ground-truth stimuli (see Wei and Stocker, 2017, for review). Importantly, Bayesian efficient coding models predict that the strength of perceptual biases is governed by the uncertainty in neural activity that encodes the stimulus (Wei & Stocker, 2015). In line with this, prior work has varied external stimulus properties (e.g. contrast, exposure time) and observed stronger biases for weaker/shorter stimuli (de Gardelle et al., 2010), putatively by manipulating the amount of neural spiking activity each stimulus received. However, the dependence of perceptual biases on neural encoding strength—a critical prediction of Bayesian models—has not been directly investigated with measurements of neural activity. Thus, there is limited insight into whether neural excitability states modulate biased perception, despite this being a major prediction of efficient coding theory.

A large portion of endogenous neural activity consists of rhythmic fluctuations that are thought to play a role in cortical information processing and communication (Fries, 2005; Jensen & Mazaheri, 2010). Alpha-band activity (8–14 Hz) is especially relevant for sensory coding as it modulates neuronal spiking in sensory cortices and covaries with visual performance (Dougherty et al., 2017; Haegens et al., 2011). The effects of alpha oscillations are primarily inhibitory, such that increased power scales inversely with neural excitability (Mathewson et al., 2011; Mazaheri & Jensen, 2010). For instance, brief visual stimuli that are presented during high alpha power are less frequently detected by observers (Chaumon & Busch, 2014; Iemi et al., 2017; Thut et al., 2006) and appear subjectively to have weaker relative contrast (Balestrieri & Busch, 2022). The phase of the alpha wave (e.g. peak vs trough) also modulates detection performance (Beale et al., 2026; Busch et al., 2009; Dugué et al., 2011; Harris et al., 2018) and the decodability of visual stimuli (Zhou et al., 2021), suggesting that alpha inhibits the fidelity of stimulus representations in the brain. These findings suggest that oscillatory states regulate the strength of sensory coding, but it remains to be seen whether this affects the precision or bias of encoded stimuli as predicted by efficient coding accounts.

While efficient coding has primarily been understood to reflect the static allocation of sensory resources (e.g. neural tuning density), recent examples suggest that repulsive biases may be actively created during online vision. In the tilt illusion, for example, the perceived orientation of a central stimulus is repulsed away from a surrounding oriented border, consistent with surround-driven changes in representational resources that may arise from a neural system adapted to centre-surround structure in natural scenes (Zhang et al., 2025). Notably, increases in

alpha power have been shown to strengthen illusory percepts in the tilt illusion (Williams et al., 2024). Given that alpha oscillations modulate neural excitability, it is plausible that oscillatory states can control the strength of repulsive perceptual biases by dynamically regulating the precision of sensory coding activity.

Here we tested the hypothesis that spontaneous fluctuations in visual excitability modulate the precision of orientation coding, and therefore the strength of repulsive cardinal biases in perception. Using closed-loop electroencephalography (EEG), we presented near-threshold gratings during periods of strong versus weak occipital alpha activity. On each trial, participants reproduced the perceived orientation and rated stimulus visibility, allowing us to quantify cardinal bias, response variability, and subjective perception as a function of ongoing alpha state. We then fitted an efficient Bayesian observer model to estimate how alpha-dependent changes in behaviour map onto latent components of sensory coding, including the allocation of representational resources (Fisher Information), sensory likelihood noise, and guessing. Consistent with our predictions, higher alpha was associated with stronger repulsive cardinal bias and greater response variability. Model-based analyses further indicated that this effect is primarily explained by increased sensory likelihood noise, accompanied by an increase in guess rate.

3.1 Results

Alpha Oscillations Produce Stronger Repulsive Cardinal Biases

We recorded brain activity using EEG while human observers viewed brief oriented gratings (randomly sampled between 0–180°) and performed an orientation reproduction task (Fig 3.1a). On each trial, an oriented probe stimulus appeared either to the left or right of central fixation. Observers then reproduced the orientation of this probe stimulus after a short delay and provided a rating of its visibility. Performance was matched across observers by setting stimulus contrast at a 'threshold' level for variability in orientation reproduction errors (see *Methods – Adaptive Psychophysical Procedure*). To assess the dependence of perceptual processing on visual alpha oscillations, we analysed online data from the average of five posterior sensors (see *Methods - Online EEG Analyses*) and presented stimuli during targeted periods where spontaneous alpha-band activity was strong or weak (Figure 3.1b). Online stimulus presentation allowed us to collect more data in high alpha power states than would otherwise be captured if relying on randomised stimulus timings to capture spontaneous alpha activity. This approach let us

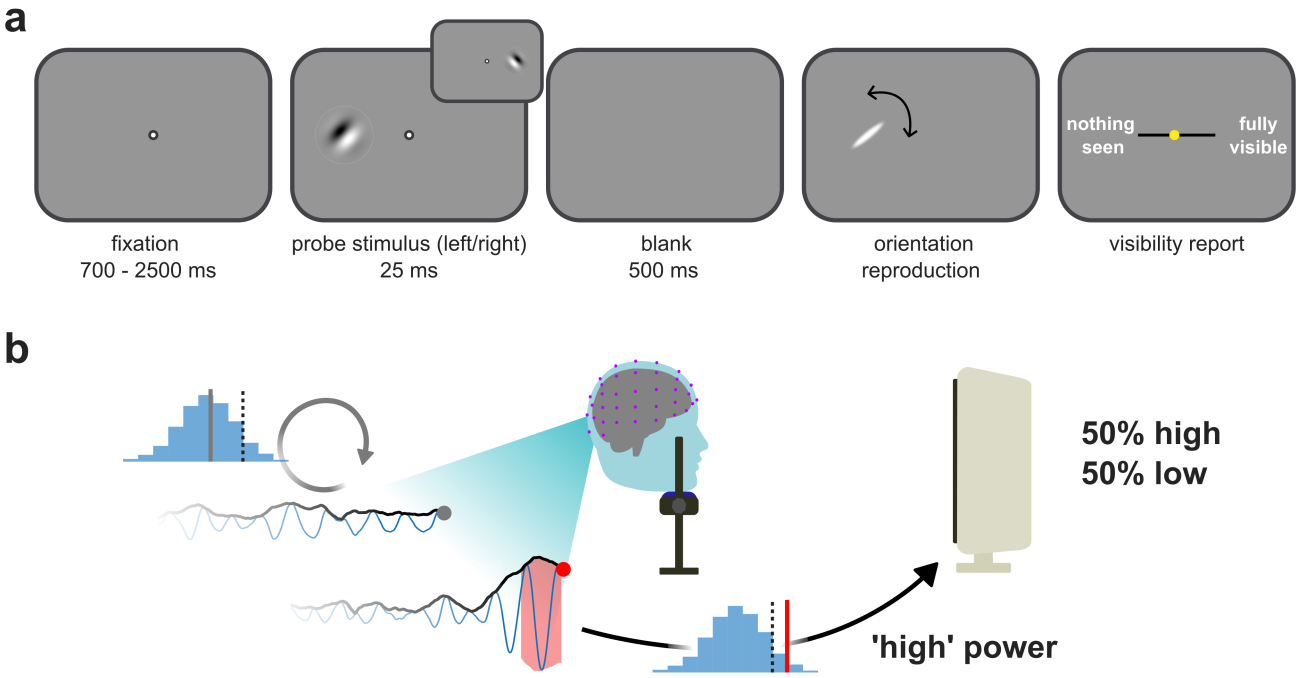

Figure 3.1: Behavioural task and online detection of alpha oscillations

a Stimulus and trial design for the orientation reproduction task. Participants viewed a brief randomly-oriented grating then gave a continuous reproduction of the stimulus and rated its visibility. **b** Online EEG analysis classified the prestimulus alpha oscillation power. Stimuli were presented adaptively during periods when alpha power was high (above 90th percentile, based on z-score) or low (below the threshold). The histograms depict the online estimate of alpha power at the latest time sample (solid grey/red line) for below threshold low power, left image, and above threshold high power, right image (compare solid line against the dashed line indicated 'high' power threshold). An adaptive algorithm maintained approximately 50% of stimulus presentations during high/low power, respectively, throughout each session.

evenly split our psychophysical analyses into a high and low power condition (see Figure 3.2a).

We used a trial-level measure of bias that was based on the response error relative to the true stimulus, of participants' orientation reports. To quantify the influence of the cardinal (horizontal and vertical) orientations, we re-signed the response errors such that positive errors indicated repulsion (i.e. the reported orientation was further from the nearest cardinal axis than the stimulus). As shown in Figure 3.2b, responses were repulsive consistent with previous reports of cardinal orientation biases (de Gardelle et al., 2010; Jastrow, 1892; Smith, 1962). Having established the presence of repulsive biases in our dataset, we used a mixed-effects General Linear Model (GLM) analysis to investigate whether repulsion differs as a function of the alpha power at stimulus presentation. This revealed

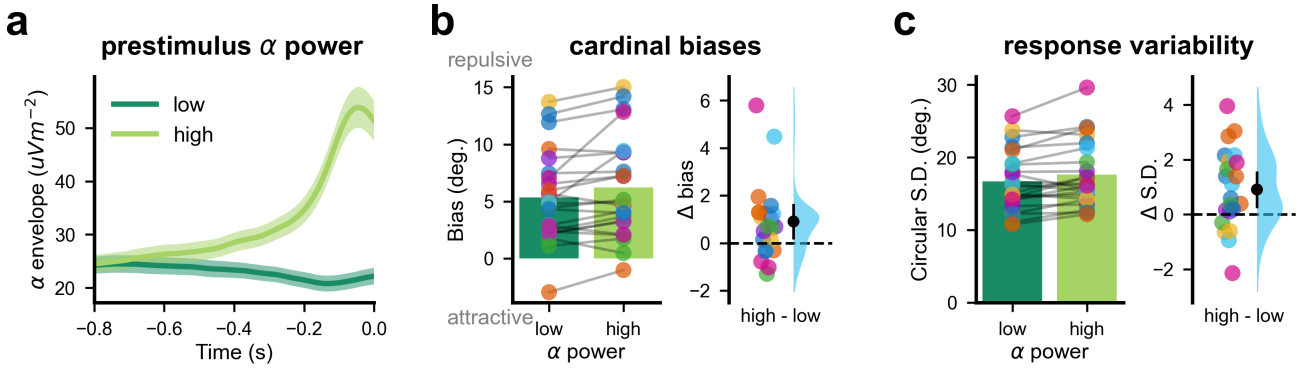

Figure 3.2: Alpha power produces stronger orientation biases and greater variability in perception

In all panels, high and low alpha power conditions are represented by light and dark lines, respectively. **a** Time series of prestimulus alpha power by condition, computed offline using an identical processing pipeline to the online analyses (values are the envelope amplitude of Hilbert transformed data). Shaded lines show the standard error of the group mean. **b** Bias of the reproduced stimulus orientations relative to the cardinal (horizontal & vertical) axes. Left panel shows the average bias for each observer (coloured dots) in high and low alpha. Positive values indicated repulsion of responses away from the cardinal axes. Light and dark bars show the group mean. Right panel shows the within-participants change in bias. **c** Orientation reproduction variability quantified using the circular standard deviation (S.D.) of responses. Conventions as in **b**.

that cardinal repulsion significantly increased in high alpha power, when neural excitability is diminished (Mathewson et al., 2011; fixed effects estimate (β) = 0.9; 95 % Confidence Interval (CI), (0.27, 1.54); $t(17755) = 2.77$; $p = .010$). We next quantified the overall variability of perceptual reports as a function of alpha power. The unsigned magnitude of trial-level response errors was entered into a GLM and we found that high alpha power predicted significantly greater response errors (β = 0.95; 95% CI, (0.38, 1.51); $t(17, 755) = 3.29$; $p = .001$). We also quantified the difference in response variability using the circular standard deviation fit to each participant's response errors ($t(22) = 3.05$; 95% CI, (0.28, 1.53); $p = .006$; Figure 3.2c).

Alpha Modulates the Noise of Efficient Sensory Likelihoods

We next investigated how alpha oscillations alter sensory representations of visual orientation and thereby shape behaviour in a stimulus-specific manner. To do this, we modelled the neural computations that underlie orientation perception as a process of efficient Bayesian inference (Wei & Stocker, 2015). In this model, the observer infers the stimulus value from noisy sensory activity that encodes orientation. A key component of this model is that the neural evidence, or *likelihood*,

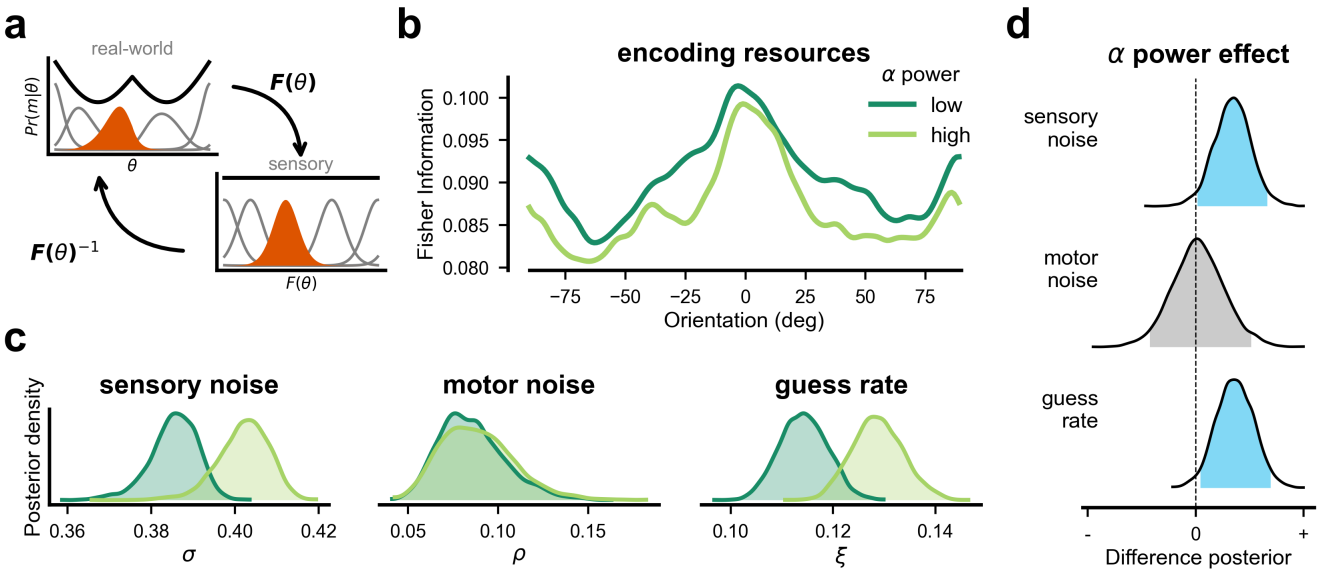

Figure 3.3: Alpha power increases likelihood noise and guesses in an efficient Bayesian observer model

a A depiction of efficient coding theory in a Bayesian inference model of perception. Stimulus likelihood distributions are symmetric in sensory coordinates (internal to the brain) and asymmetric when expressed in real-world orientation coordinates. Black line shows the allocation of neural encoding resources over real-world orientation. The function F captures the transformation of orientation into the sensory-space, where the encoding resources are uniform and the likelihood is symmetric. **b** Encoding resources fitted to orientation reports in the high and low alpha power conditions (light and dark green lines). **c** Posterior parameter probability distributions of three model parameters, from MCMC sample using kernel-density estimation. Sensory noise describes the width of the Gaussian likelihoods in the efficient sensory space. Motor noise and guess rate account for the variability in orientation reproduction that is not captured by the efficient Bayesian inference. **d** Posterior distributions of the parameter differences in normalised parameter units. The shaded region of the distribution is the 95 % highest density interval (HDI) and is shown in blue when the HDI does not contain zero.

of the stimulus is shaped by how encoding resources are non-uniformly allocated over the stimulus domain. If the brain allocates more of its encoding resources to cardinal orientations (i.e. under the efficient coding principle), then the model predicts repulsive cardinal biases. This arises from the fact that sensory noise in the efficient representational space does not map uniformly onto the external world, producing sensory likelihoods that are skewed away from densely encoded regions of the stimulus space (see Figure 3.3a, e.g. Wei and Stocker, 2015).

We followed the recent modelling approach of Hahn and Wei (2024) in fitting the allocation of stimulus encoding resources as a free parameter of our model. This let us quantify observers' representational capacities directly from

behavioural responses. We fitted the efficient Bayesian observer model to the continuous trial-level orientation judgements, separately for high and low alpha power conditions. Figure 3.3b shows the shape of the encoding resources as a function of alpha power, in units of Fisher Information (FI) over stimulus orientation. Encoding resources peak at cardinal orientations as expected under efficient coding. This feature is present for stimuli presented during both high and low alpha power. Note that we examined the FI qualitatively, rather than statistically, as it is a latent variable that redescribes the patterns of behavioural bias and variance reported above.

We next examined how alpha power affects further parameters of the model that shape the translation of internal representation to behavioural response. Figure 3.3c shows the posterior parameter distributions of three model parameters that describe the sensory (likelihood) noise, motor noise, and guess rate. We assessed the effect of alpha oscillations on these parameters using the posterior difference distributions between high and low alpha (see Figure 3.3d). We found that high alpha power produced an increase in the sensory noise parameter, which corresponds to the width of a Gaussian stimulus likelihood inside the efficient sensory space (posterior difference mean (M) = .017; 95% Highest Density Interval (HDI), (0.001, 0.034)). The motor noise parameter captures additional variance in orientation reports that is not biased by encoding resource allocation, and we did not observer changes in motor noise due to alpha power (M = 0.003; 95% HDI, (-0.048, 0.062)). We also found that high alpha power produced a credibly larger guess rate (M = 0.015; 95% HDI, (0.001, 0.028)), corresponding to a greater proportion of trials that were effectively random responses with uniform orientation. Thus, our efficient Bayesian observer modelling suggests that alpha power affects orientation reproduction behaviour via two mechanisms: (1) increased sensory likelihood noise, and (2) a higher rate of guesses. While both mechanisms increase the variance of orientation reports, only the increased likelihood noise produces stronger repulsive biases.

Our previous analysis of the re-signed response errors provided a simple quantification of cardinal biases that allowed us to collapse across stimulus orientation. We plot the profile of stimulus-specific response errors in Figure 3.4a. We observed a strong repulsive cardinal bias that is consistent with our above analysis and with previous studies (de Gardelle et al., 2010). Our efficient Bayesian observer model also produced cardinal biases that closely matched behaviour (Figure 3.4b). Figure 3.4a highlights that the strength of repulsive biases was strongest around the

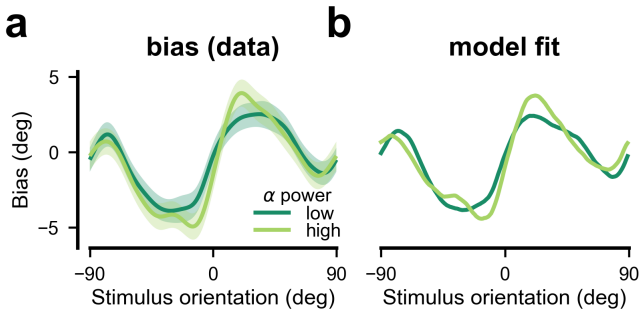

Figure 3.4: Stimulus-specific biases in reported orientations

a Bias in behavioural data estimated using a sliding Gaussian window (10° standard deviation). The positive slope at zero and ±90 reflects a repulsive bias away from cardinal orientations. Light and dark lines show high and low alpha power, respectively. Shaded regions show the 95% confidence interval estimated using bootstrapping. **b** Bias in the orientation inferences of the efficient Bayesian observer model closely reproduce participant data.

vertical axis (0 degrees) in our data. This unexpected finding held in the re-signed errors; we found less repulsion around horizontal orientations relative to vertical (β = -2.68; 95% CI, (-3.49, -1.87); $t(17,755)$ = -6.47; $p < .001$) and a significant interaction with alpha power (β = -1.48; 95% CI, (-2.63, -0.33); $t(17,755)$ = -2.52; p = .012) indicating that alpha-related repulsion was significant around the vertical axis (β = 1.60; 95% CI, (0.79, 2.41); $t(8925)$ = 3.89; $p < .001$) but not around the horizontal axis (β = 0.12; 95% CI, (-0.67, 0.91); $t(8828)$ = 0.31; p = .760).

Visibility is Modulated by Alpha Oscillations

We investigated how alpha oscillations affected the visibility of oriented stimuli. Participants subjectively rated the visibility of each stimulus on a continuous scale (see *Methods*). While individuals differed in their usage of the scale (Figure 3.5a), we used a mixed-effects GLM to predict relative changes in visibility as a function the alpha power state. Stimuli presented during high alpha power were rated as less visible than those presented during low power (β = -1.48; 95% CI, (-2.34, -0.62); $t(17,755)$ = -3.38; $p < .001$). Figure 3.5b shows the stimulus-specific visibility ratings that were estimated using a sliding Gaussian window, which shows the largest decrease in visibility for vertically oriented stimuli where the repulsive bias is strongest in our dataset. The instantaneous phase of prestimulus alpha oscillations was entered into a mixed-effects GLM using the sin and cos of phase values as predictors. This revealed that visibility was also predicted by the specific phase of

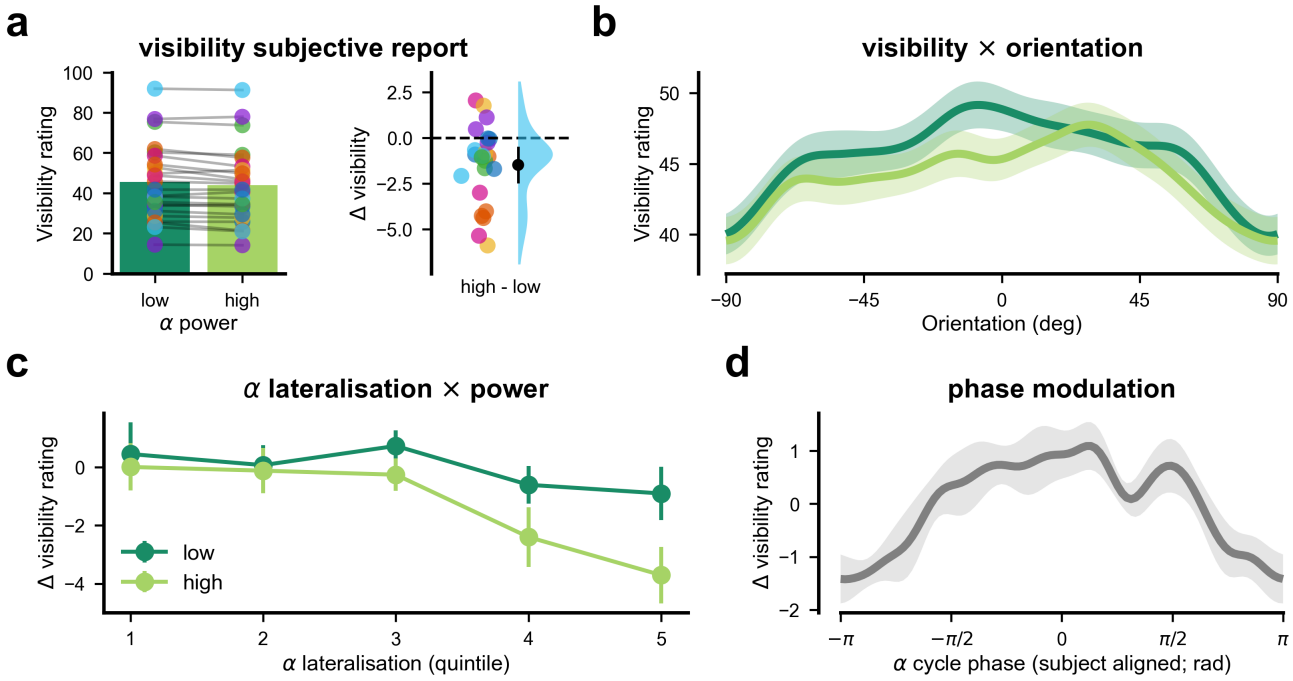

Figure 3.5: Alpha oscillations reduce visibility ratings

a Visibility rating, split by alpha power condition, from a continuous subjective rating scale (see task in Figure 3.1a). Left panel shows visibility ratings on the original scale units (0–100). Right panel plots the within-participants change in visibility between high and low alpha power conditions. **b** Stimulus-specific visibility ratings estimated using a sliding Von Mises window over stimulus orientation. Shaded lines show the 95% confidence interval based on bootstrapping participant-level data. **c** Prestimulus hemispheric alpha power lateralisation and its effect on within-participants changes in visibility (within participants standard errors; Morey, 2008). The x-axis shows a quintile split of lateralisation. Positive numbers indicate that stimulus-contralateral hemispheric alpha power was greater than stimulus-ipsilateral. **d** Alpha oscillation phase effect on within-participants visibility ratings. Visibility ratings were averaged over prestimulus phase using a sliding circular-Gaussian window. The x-axis shows the oscillatory phase of alpha aligned by the subject specific phase with highest visibility (estimated from mixed-effects GLM analysis using the sin and cos of phase as predictors).

alpha oscillations at the time of stimulus presentation (Figure 3.5d; $F(2, 17752) = 6.26$; $p = .002$).

As the stimuli were presented in lateral locations, we were able to examine how visibility depended on alpha power in the specific hemisphere of the brain that processed the stimulus. To quantify hemispheric lateralisation, we computed offline the power in left- and right-hemisphere electrodes and subtracted the stimulus-contralateral power from the stimulus-ipsilateral power. Figure 3.5c shows the change in visibility using a quintile split of alpha lateralisation, additionally split

by high or low power at stimulus presentation (quintile binning was used for visualisation while analyses were conducted on continuous data). We found that hemispheric lateralisation alone did not predict changes in visibility (β = -0.09; 95% CI, (-0.83, 0.64); $t(17, 753)$ = -0.25; p = .805) but did interact with the alpha power at stimulus presentation (β = -0.98; 95% CI, (-1.83, -0.13); $t(17, 753)$ = -2.27; p = .023), such that decreased visibility depended on there being high alpha power in the hemisphere that processed stimulus.

We performed similar analyses that used alpha phase and power lateralisation to predict repulsive biases and variability of orientation reports. These analyses did not reveal any effect of either phase or lateralisation on these behavioural measures (all ps > .05).

3.2 Discussion

Horizontal and vertical orientations are overrepresented in the natural environment (Girshick et al., 2011) and, in an efficient neural system, this leads to perception that can be repulsed away from these cardinal orientations (Wei & Stocker, 2017). We investigated whether the perceptual predictions of efficient coding are modulated online by endogenous neural excitability. In particular, we examined whether occipital alpha oscillations influence the repulsive cardinal bias inherent to orientation perception, and the subjective visibility of oriented stimuli. We demonstrated that stimuli presented during strong alpha activity are reported as both further biased away from cardinal orientations and less subjectively visible. To explain this, we used a Bayesian model of perceptual inference that efficiently allocates sensory encoding resources over orientation stimuli. The model adopts the constraints of realistic sensory systems, which are tasked with efficiently encoding stimuli that have non-uniform prior distributions in the environment. The model suggested that stronger repulsive biases during higher amplitude alpha oscillations come from increased noise in the sensory representations of visual orientation.

Visual orientation is encoded in the activity of orientation-selective cortical neurons (Hubel & Wiesel, 1962; Vogels, 1990). A popular idea is that neural population activity approximates a Bayesian 'likelihood' function that allows external stimuli to be read out probabilistically from the neural code (Jazayeri & Movshon, 2006; Ma et al., 2006; Pouget et al., 2000). Uncertainty in the stimulus likelihood may then arise when the neural code is corrupted via internal noise or weakened neural activation, as both of these result in fewer spikes to effectively encode the

stimulus. For example, one study manipulated the duration of oriented stimuli to produce weaker processing with relatively shorter presentations (de Gardelle et al., 2010). This manipulation produced repulsive cardinal biases that scaled inversely with stimulus duration, consistent with the above account. Additionally, repulsive orientation biases are stronger when stimuli are presented in larger set sizes, as they may receive weaker spiking activity due to inhibition from divisive normalisation (Taylor & Bays, 2018). Our results show that alpha oscillations, which inversely correlate with neural excitation, correspond to changes in perceptual biases, supporting the putative link between neural excitability and visual perception. We suggest that alpha-band oscillatory activity can impact the quality of stimulus codes in early visual neurons, consistent with its inhibitory effect on visual spiking (Dougherty et al., 2017; Haegens et al., 2011). The resulting noise variation in perceptual estimates could reflect reduced signal-to-noise levels from suppressed spiking activation across the neural populations that encode orientation.

The inhibitory effect of alpha oscillations on neural activity has been shown via direct neural recordings (Dougherty et al., 2017; Haegens et al., 2011) and underpins theoretical models of oscillatory organisation of cortical processing (Fries, 2005; Schalk, 2015). A range of behavioural experiments suggest that alpha alters perceptual performance, but it has been difficult to interpret the precise effect of alpha on sensory representational ability. For instance, stimuli presented during spontaneous increases in alpha power are less likely to be detected by humans, although fewer 'present' reports were also found in a stimulus-absent condition (Limbach & Corballis, 2016). Furthermore, alpha power does not seem to affect performance in forced choice discrimination tasks (Iemi et al., 2017), but does decrease confidence in such tasks (Samaha et al., 2017a). These findings suggested that alpha oscillations may not affect sensory representations in the context of perceptual discrimination, but rather impart an additive shift in general neural activity that can improve meta-cognitive evaluation and detection when alpha power is reduced (Pilipenko & Samaha, 2024; Samaha et al., 2020). Our findings recast this and show a previously unrealised effect of alpha power on the uncertainty within fine-grained sensory representations.

We found that the feature-specific content of perceptual representations is altered by alpha oscillations, and that this produces more distorted and more variable judgements of orientation. A relevant difference is that past work has probed discrimination ability using forced choice judgements across distinct stimulus categories (e.g., Gabors rotated 10 degrees clockwise or counterclockwise

from vertical; Iemi et al., 2017). It may be that the neural representations that are relevant to performance in those tasks are largely categorical, which differs from the fine-grained stimulus reproductions that our observers were required to generate. An interesting exception is that alpha power affected categorical discrimination about stimuli based on the tilt illusion, such that the illusion is more consistent in high alpha power (Williams et al., 2024). Given that the tilt illusion arises from efficient coding mechanisms that enhance neural coding of the surround orientation (Zhang et al., 2025), stronger neural inhibition in high alpha is expected to increase neural uncertainty, just as we observed here, and so should strengthen repulsive biases relative to the surround. Thus, our findings suggest that the neural representation of visual features is weakened under strong alpha power, producing perceptual estimates to be repulsed from stimulus regions that have stronger coding fidelity.

We measured subjective reports of stimulus visibility and found that high alpha power led to weakened visibility ratings. Our modelling also showed that strong alpha power produced a higher guess rate. This suggests that weakened stimulus processing also impacted objective perceptual behaviour; specifically, alpha power increased the likelihood of reporting a uniform random orientation unrelated to the stimulus. Furthermore, we show that alpha phase predicts visibility ratings on a trial-to-trial level (c.f. Benwell et al., 2017, 2022; see also Harris, 2023). However, our analyses did not detect a similar relationship between phase and perceptual bias, counter to existing evidence that perceived visibility predicts orientation biases even more strongly than physical manipulation of stimulus duration (de Gardelle et al., 2010). Thus, alpha oscillations may affect the noise of visual coding which can bias downstream perceptual decoding and behaviour, and shape meta-cognitive evaluations.

In summary, here we have shown that the representational fidelity of efficient sensory codes are shaped by endogenous oscillatory activity in the alpha band. We thus demonstrate that oscillating inhibition of cortical activity has measurable influences on behavioural biases and variability in human visual perception. What is more, these effects on perceptual inference are also apparent to observers, as reflected in their subjective ratings of stimulus visibility. Overall, our results underscore that perceptual inferences are dynamically shaped by long-term influences, like environmental image statistics, and ongoing neural excitability states.

3.3 Methods

Participants

Thirty-one healthy adult humans were recruited into the study. The study was approved by The University of Queensland Human Research Ethics Committee. Participants gave their written informed consent and were paid $20/hour. Our sample size was not based on power calculations; rather, we aimed to test as many participants as feasible given constraints on project funding and resources. We required participants to complete an initial screening session consisting of 400 trials of the psychophysical task without EEG, to develop their familiarity and comfort with the task. One person was removed at this stage for having poor behavioural performance and low engagement with the task. A further six participants were removed from the sample due to self-attrition at the follow-up EEG session, and another person was removed due to problems with the EEG recording. The final sample consisted of 23 participants with mean age of 23 years (19–30 years old, 18 females). All participants had normal or corrected-to-normal vision.

Apparatus

Visual stimuli were displayed on a 24-inch ViewPIXX View3d monitor (1920x1080 pixels; 120 Hz refresh rate) placed in a dark, electromagnetically shielded room. Psychtoolbox (Brainard, 1997) v.3.0.17 was used to control stimulus presentation in MATLAB (MathWorks, R2020a). Viewing distance was maintained at 57 cm with a chin-rest. The screen was gamma-corrected using an X-Rite I1 colorimeter. A Biosemi ActiveTwo system (Biosemi, Amsterdam, Netherlands) was used to record EEG from 64 Ag-Cl electrodes arranged in a standard 10-10 layout (Oostenveld & Praamstra, 2001). Lab streaming layer (Kothe et al., 2025) was used on an additional PC (Dell) to read EEG data into MATLAB for online processing whilst also saving to hard disk. The realtime state of alpha oscillations was conveyed to the stimulus computer via parallel port connection and read into MATLAB to control Psychtoolbox presentation.

Stimuli and Task

Visual orientation was probed using sinewave gratings (1.5 cycles per degree of visual angle [dva]; random phase) presented within a Gaussian spatial window

(0.33 dva sigma) on a mid-grey background (Figure 3.1a). The probe stimulus was randomly positioned on each trial to either the left or right side of fixation with 6.5 dva displacement along the horizontal meridian. A low contrast circle (1 pixel wide; 3 sigmas radius) surrounded the stimulus to reduce spatial uncertainty upon its appearance. Central eye fixation was maintained during the prestimulus period with a small black dot (0.3 dva diameter) containing a white inner dot. The response stimulus was an elongated white bar centred on the probe location.

Each trial started with a variable fixation period of 700–2500 ms followed by a randomly oriented probe stimulus lasting 25 ms. The observer reproduced the presented orientation by adjusting a response stimulus with a computer mouse. Cursor position was indicated with a low-contrast grey dot (0.33 dva diameter) and was visible at the centre of the probe location following a blank 500 ms period after the stimulus. To minimise bias from random setting of the response stimulus orientation, the response stimulus was hidden from view until the cursor had moved more than 0.3 dva from the start position, and then reflected the direction in which the cursor had been moved. Reproduction responses were untimed and were selected using a mouse button to record the target stimulus orientation. Observers then rated the subjective visibility of the probe stimulus using a sliding rating scale with the onscreen text anchors 'nothing seen' to 'fully visible'. The EEG task consisted of 800 total trials, with voluntary breaks given after every 4 minutes on task. Feedback was given on the first 30 trials (in both behavioural and EEG sessions) by overlaying the true probe orientation using two green dots placed at the outer circumference of the imaginary ring containing the response stimulus. These feedback trials were removed from analyses.

Adaptive Psychophysical Procedure

Performance in the orientation reproduction task was maintained at a criterion level of difficulty using an adaptive psychophysical procedure to set the contrast of probe stimuli. Based on pilot data, we confirmed that the precision, κ, (or inverse variance) of orientation reports is well-described by a function over stimulus contrast, c,

$$\kappa = \kappa_{max} \frac{c^m}{c^m + c_{50}^m} \tag{3.1}$$

such that response precision is zero at $c = 0$ and rises non-linearly with slope, m, to a maximum level κ_{max}. The parameter c_{50} locates the midpoint of the slope and is the contrast value at which precision changes most rapidly. For this reason,

c_{50} is highly informative for understanding the contribution of neural variability to behavioural performance. Function parameters were fit online using a custom particle filter implementation (DiMattina, 2015). An adaptive block of 80 trials was presented at the start of the session, in which stimulus contrast was set to the value that maximised the expected information gain of the Bayesian posterior over parameter values. During the main experimental block, stimulus contrast was presented at the up-to-date estimate of c_{50} and the particle filter was updated on each trial, to ensure presentation using the most-current estimate of the model parameters.

Electroencephalography

Online Processing and Presentation

EEG data were read into MATLAB in real-time and preprocessed continually, using a four-second epoch that was appended with the latest available samples. Sensors were average-referenced, demeaned, and filtered with a second-order Butterworth bandpass between 1–100 Hz. We then applied a surface Laplacian transform using precomputed matrices. Signals were averaged in five posterior electrodes (Oz, O1, O2, PO7, PO8) to create a posterior region-of-interest. Following an improved method for estimating bandpass activity (detailed in Schreglmann et al., 2021), we used a causal second-order Butterworth filter to apply a bandpass (8–14 Hz) in the frequency domain and return the Hilbert-transformed analytic signal time series data. This method, termed the endpoint-corrected Hilbert transform, has been shown to provide estimates of instantaneous phase and amplitude that are not biased by artifacts from short time-series (Schreglmann et al., 2021).

We then used the absolute value of the analytic signal as a measure of alpha power, which was averaged in the latest 50 ms of online data, log-transformed, and z-scored. A rolling estimate of the signal mean and variance was tracked only when the stimulus PC indicated that the participant was in the prestimulus fixation period. Alpha power was classed as high when z-score amplitude exceeded the 90th percentile of the normal distribution. A separate stimulus PC received the realtime alpha state during the fixation period (e.g. 1 for high alpha, 0 for not high alpha) and used this in an adaptive algorithm that sought to present a stimulus in high power on ~50% of trials over the experiment. This was done by either waiting for a high alpha signal within the fixation period (if less than 50% of past trials were classed as high alpha presentations) or setting a randomised fixation interval before the stimulus was presented. In testing, we observed an eighty-three millisecond

delay in online EEG processing, meaning alpha-band changes could be captured in approximately one cycle of prominent activity.

Offline Processing

The EEGLAB v2024.2 toolbox was used to perform offline processing of EEG. Signals were down-sampled to 512 Hz and high-pass filtered at 1 Hz using *pop_eegfiltnew.m* with default filter parameters for the purpose of data cleaning. Data was then epoched between -1.5–2 s centred on stimulus onset and was average-referenced. Bad channels were interpolated via spherical splines and automatically detected by adaptively thresholding (at 3.5 $|z|$) the per-channel variance and average across channel correlations. Laplacian filtering was applied as above. We also adaptively rejected trial epochs that exceeded more than 3 $|z|$ in signal range, deviation from channel grand average, and variance over the epoch (average of channels). Rejected trials were removed from all analyses (average rejected, 2.75% per participant; range, 0.7–5.9%). As we recorded online alpha amplitude on each trial but not phase, we averaged channels in the ROI used for online processing and applied the same bandpass filtering to recover the endpoint-corrected alpha phase at stimulus onset, as this was not stored online due to the latencies in online processing. Alpha power lateralisation was computed by bandpass filtering data averaged within a left (O1, PO7) and right (O2, PO8) region-of-interest and subtracting the endpoint-corrected amplitude at stimulus onset (stimulus-contralateral minus stimulus-ipsilateral), and z-scoring within each participant.

General Linear Modelling

We quantified cardinal biases on a single-trial level using re-signed response errors. This was computed by circularly subtracting the stimulus orientation from the response, and then resigning errors based on whether the response was further or closer (positive or negative, respectively) to the nearest cardinal axis than the stimulus was. This yields a non-zero-mean distribution of errors if there is cardinal repulsion or attraction across trial-level responses. Response variability was quantified on a single-trial level using the absolute value of response errors as measure of error magnitude.

We fit mixed-effects General Linear Models (GLMs) to these data and included random slopes and intercepts in the models. Separate models were used to predict re-signed errors, absolute errors, and visibility reports (dependent variables), using the phase and amplitude of alpha oscillations as predictors. Models were fit with

MATLAB's *fitglme.m* function. The resulting fixed-effects β coefficients predict the change in the outcome variable expressed in degrees of orientation, where relevant. For phase analyses, we included predictors for the sin and cosine of prestimulus phase values. The mixed-effects model allowed us to infer the optimal phase ϕ of each individual j using the fitted random effects, and reconstructing via $\phi_j = \text{atan2}(\sin_j, \cos_j)$.

Bayesian Observer Model

We use a Bayesian observer model that was described by Hahn and Wei (2024), and we refer readers to this work for extensive details on implementation. Here, we give the conceptual foundations of this model and describe adaptations that let us fit the model using Monte Carlo inference. The core idea is that stimuli θ are encoded into sensory measurements m that are corrupted by noise inside representational space

$$m = F(\theta) + \delta \tag{3.2}$$

where δ is additive sensory noise (von Mises). The function F maps stimuli into sensory space (i.e. the encoding) and thus determines how much of the limited representational space is allocated over certain values of θ. The shape of this allocation is formally described by the Fisher Information $J(\theta)$, which is proportional to the squared rate of change in the encoding function F:

$$J(\theta) = \frac{F'(\theta)^2}{\sigma^2} \tag{3.3}$$

where σ^2 is the sensory noise variance. The observer model uses Bayes' rule to obtain a posterior distribution $p(\theta|m)$ given the sensory measurement likelihood $p(m|\theta)$ and stimulus prior distribution $p(\theta)$. For each sensory measurement, the observer model applies Bayesian decision theory to select a point estimate of the stimulus $\hat{\theta}$ that has the least expected loss. Behavioural responses are modelled as the value of $\hat{\theta}$ that is corrupted by additional von Mises motor noise ρ and the probability of uniform response ξ, which accounts for the influence of guessing on some trials.

To implement the observer model, we mapped stimulus values to the closest position on a discrete grid of $N = 180$ values and calculate $p(\theta|m)$ via exact Bayesian inference, projecting sensory measurement likelihoods onto grid points. We used a uniform stimulus prior and a cosine L_p loss function as in Hahn and Wei (2024) with $p = 8$ (which was the best fitting model of $p \in \{1, 2, 4, 6, 8\}$,

in agreement with Hahn & Wei's results). Resource allocation (i.e. the stimulus likelihood) was fitted non-parametrically by using a free parameter for $F'(\theta_n)$ at each grid point n. A regularisation term was used to penalise the squared difference between adjacent grid points, and we treated this as an additional parameter during model inference. We used Hamiltonian Monte Carlo sampling to perform Bayesian inference of the observer model. In doing so, we diverged from the implementation in Hahn and Wei (2024) when computing the Bayes estimator $\hat{\theta} = \arg_{\hat{\theta}} \min \int L_p(\theta, \hat{\theta})\, p(\theta|m) d\theta$. We instead computed the expected loss at each discretised grid point and chose the grid value with the lowest loss as the Bayes estimate $\hat{\theta}_n$.

Sliding Window Analyses

We used a sliding window approach to quantify stimulus-specific variation in orientation reproduction bias and visibility ratings. Data were averaged within a circular normal (von Mises) window with a standard deviation of 5°. Bias was computed by estimating the circular mean of orientation responses and then circularly subtracting the centre of the window. Bootstrap resampling was used to estimate the variability in sliding window estimates. We also used the sliding window approach to visualise the phase modulation of visibility (Figure 3.5d; 15° standard deviation).

Statistical Analyses

The Bayesian observer model was constructed in PyMC v5 (Abril-Pla et al., 2023) and fitted using the No U-Turn Sampling method for Hamiltonian Monte Carlo with GPU acceleration via JAX. $\hat{R}$ values were less than 1.01 for all parameters, suggesting that sampling chains converged well to the target posterior. For statistical inference we analysed the posterior distribution of parameter differences between high and low alpha power. This was computed by subtracting the high from low power condition, per sample. The 95% Highest Density Interval was used to infer whether the difference distributions credibly lie outside the zero-value (i.e. the parameter difference expected under the null). In the GLM analyses, we assessed significance using the p values of fixed-effects coefficients. To test phase effects we used the *coefTest.m* function to conduct a linear hypothesis test that both the fixed-effects of the sin and cosine predictors were zero.

Chapter 4

A Neural Observer Model of Rhythmic Perception

Understanding how neural oscillations shape behaviour needs both empirical and theoretical work. Empirical studies can establish reliable associations between oscillatory states and perceptual choice, but those associations remain incomplete without a clear account of the computations that link ongoing neural dynamics to behaviour. The aim of this chapter is to develop a mechanistic observer model that links alpha-band inhibition to sensory coding and, through that coding, to perceptual decisions. This provides a formal framework for interpreting the experimental findings from the previous chapters and for generating predictions that extend beyond the specific tasks studied there.

The core ideas behind the model are simple and threefold: (1) Visual neurons encode information about the stimulus in their firing rates; (2) alpha oscillations modulate those rates through rhythmic inhibition; and (3) decisions depend on how downstream systems read out the resulting population activity. This chapter explains each of these core ideas, which have a basis in existing and often longstanding theory of neural computation, before applying it to the empirical data from the previous chapters. I first outline a standard model of visual population coding based on receptive fields, divisive normalisation, and stochastic spike generation. I then use Bayesian decision theory to describe how the information carried by those spikes can support detection and estimation. Finally, I formalise the pulsed-inhibition account of alpha activity and show how it can be incorporated into the encoding model.

Bayesian decision theory is used here as a normative bridge between neural representation and behaviour (Yuille & Bülthoff, 1996). I do not claim that the brain

literally computes exact posteriors from a probabilistic representation. Rather, using Bayesian decision theory provides a useful way to describe the information that is available in the neural code and how that information should be used under ideal assumptions. The modelling introduced in this chapter is therefore an 'ideal observer model' (Geisler, 2011) that describes the best achievable performance in a perceptual task, based on the stimulus characteristics and their processing under varying alpha-band states. Note that, in this context, optimality is a property of the model rather than an assumption about how humans behave, so any divergence between model predictions and actual behaviour is allowable and theoretically informative.

The end goal of this chapter is to apply the model to the contrast detection task described in Chapter 2 and the orientation estimation task used in Chapter 3. The results produced by these tasks provide a way to test claims about the specific perceptual mechanisms that are influenced by alpha-band oscillations. The core features of my empirical findings are captured by the computational model, which is shown through both simulations and model fitting. Following this, I note some interesting predictions of the model that provide directions for future empirical work with a strong basis in the theory developed here. The core contribution of this model lies not just in its ability to explain existing findings, but also in the fact that it makes meaningful predictions about the ways perceptual experience should co-vary with alpha-band oscillations.

4.1 Computation in Visual Neurons

The task of the visual system is to make sense of the outside world using image data received at the retina. Neuronal activity in the visual system must contain sufficient information about the stimulus to convey its relevant features to other brain systems that serve cognition and action. The purpose of this section is to specify a compact encoding model from which population activity, uncertainty, and task-relevant evidence can be derived. I will outline the core computations needed to understand how early visual responses encode sensory information into spiking activity, which is communicated downstream as a firing rate.

Linear Filtering and Exponentiation

The receptive field is at the core of many models that seek to understand the neural code used by the early visual system. Neurons inherit their receptive fields through

the inputs they receive from earlier neurons and the way that those inputs relate to properties of the stimulus. One way to describe this is that neurons perform a simple linear weighted summation of their inputs, also known as linear filtering. Linear filtering occurs at many stages of the early visual system (Carandini et al., 2005). For instance, the centre–surround receptive fields of neurons in the retina and LGN (Enroth-Cugell & Robson, 1966) arise from linear filtering of signals from photoreceptors and retinal ganglion cells, respectively; that is, by taking the difference between two inputs that have Gaussian receptive fields corresponding to the surround and the centre. A similar mechanism accounts for the orientation tuning of V1 cells, whereby V1 receptive fields are spatially elongated because they sum synaptic inputs from LGN cells whose receptive fields lie along one spatial axis (Hubel & Wiesel, 1962).

Linear filtering describes the feature selectivity of receptive fields, such that the output of linear filtering is greatest for images that are identical to the receptive field. However, biological neurons output spiking activity that is not linearly related to the input-drive to the filter: spiking responses show threshold-like behaviour (becoming active only once a certain input strength is exceeded) and can saturate at high levels of input-drive. This suggests that linear filtering alone will not explain the full input-output transducer function of visual neurons. To address this, many models chain the linear combination of neural inputs with a fixed, or static, nonlinear operation. This forms the linear-nonlinear (LN) model, which describes real neurons well and is a good approximation of more complicated models of the biophysical dynamics behind neural spiking (Ostojic & Brunel, 2011).

A basic LN model can be captured by the following equation, describing the firing rate, λ, of a single neuron in a population of D total neurons:

$$\lambda_d(s) = (\sum_j w_j s_j)^\eta, \qquad d = 1, \ldots, D. \tag{4.1}$$

Here, exponentiation (η) is used as a simple static nonlinearity[1] that shapes the gain of each neuron's response strength. The inner term denotes the weighted sum (i.e. linear filter), where w_j is the weight applied to the jth input, s_j, for each

[1]Other nonlinear functions can also be used, and it is common to include half-wave rectification to account for the fact that spike outputs are strictly positive (Heeger, 1992).

of the neuron's J inputs, $j = 1 \ldots J$. The inputs, s_j, could be the spikes of the presynaptic neurons or the direct image characteristics of the stimulus[2].

A useful abstraction is to consider that the weighted summation in Equation 4.1 can be expressed by a neuron's *tuning function*, $f_d(s)$:

$$\sum_j w_j s_j \approx f_d(s), \qquad f_d(s) = s_c e^{\kappa \cos(s_\theta - s_d^*) - \kappa}. \tag{4.2}$$

The tuning function models the neuron's input-output transformation as a function of the stimulus characteristics directly, instead of signals from afferent neurons in the visual pathway. This is a useful move because it abstracts away details of how the stimulus is transduced in the retina and the sequence of action potentials within intermediary neurons between the retina and V1. Hubel and Wiesel (1962) showed that this move is empirically justified for the case of stimulus orientation. They found that neurons in V1 have receptive fields sensitive to oriented spatial contrast changes, such that neural responses are proportional to how well the stimulus matches a neuron's preferred orientation stimulus.

The tuning function used in Equation 4.2 takes, as input, the external stimulus s and returns a normalised response value between 0 and 1. Moving forward, the stimulus s may also be vector-valued, such that it describes both the contrast, s_c, and orientation, s_θ, value of a stimulus. Each neuron in the population has a preferred, or optimal, stimulus orientation s_d^* which elicits its strongest responses. The parameter κ controls the selectivity, or tuning width, of neural responses around those preferred orientations. A convenient assumption is that the neurons share the same average tuning width, $\kappa_d = \kappa$, and that neurons evenly tile the stimulus space (see example in Figure 4.1a). This makes the modelled population homogeneous in terms of its representational coverage (i.e. all stimulus values are represented equally), which is a property that will be exploited later in the section on efficient coding of orientation.

Divisive Normalisation

The LN model, by itself, cannot explain many of the nonlinear response properties of visual neurons. In particular, visual neurons exhibit saturating growth in their response strength when driven by a stimulus (Albrecht & Hamilton, 1982) and

[2]'Reverse correlation' is a relevant method that infers how neurons preferentially weight different features of a known sensory input. This is typically how receptive fields are determined experimentally (see Ringach, 2004, for review).

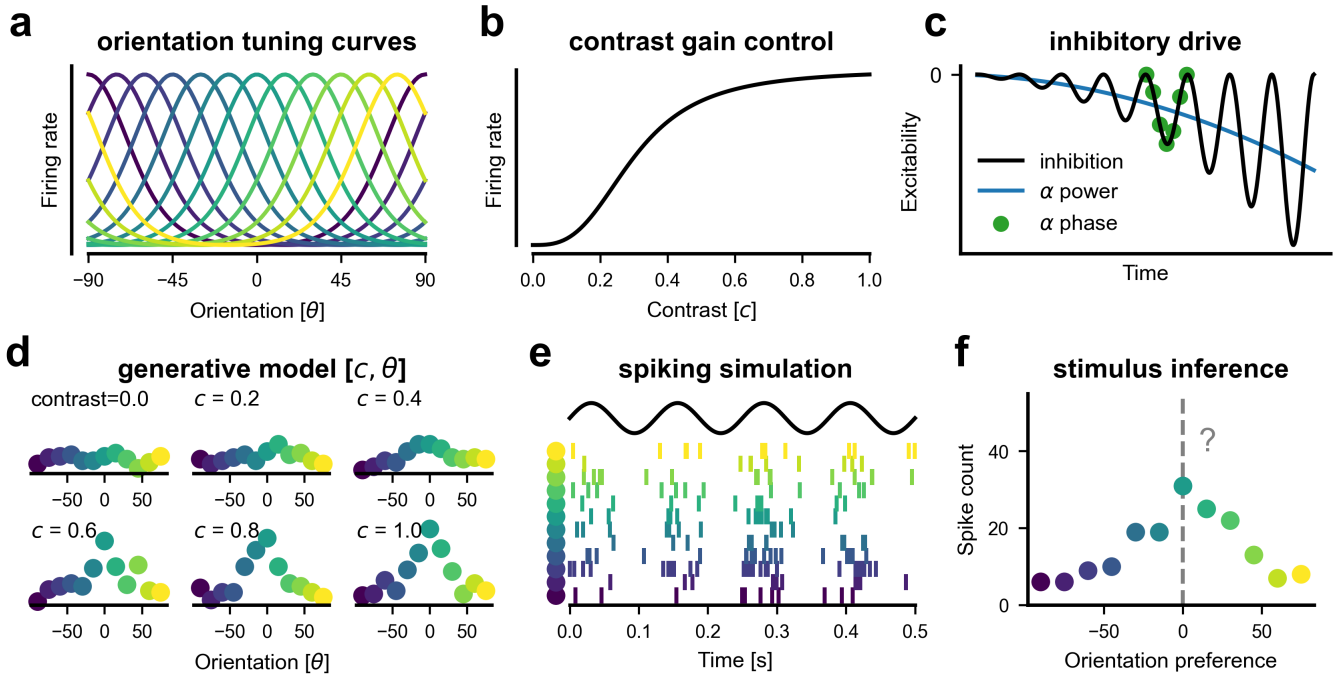

Figure 4.1: A spiking neural code model with orientation tuning, contrast gain control, and pulsed-inhibition

A neural observer model explains how a population of visual neurons encode stimuli using a combination of orientation tuning curves (**a**) and contrast gain control (**b**). **c** The pulsed-inhibition model shows how alpha power and phase are a product of inhibitory pulses with a 100 ms period corresponding to the alpha frequency. **d** Spikes can be simulated as a function of stimulus contrast and orientation. **e** Pulsed-inhibition (sinusoidal black line) influences spike generation, suppressing firing rates when inhibition is strong. **f** The influence of pulsed-inhibition on perceptual behaviour can be investigated by decoding the spiking activity for information that informs a perceptual task like detection or orientation estimation.

their responses when multiple stimuli are present cannot be (linearly) predicted from their responses to a single stimulus. Divisive normalisation was proposed to address these properties by implementing a form of nonlinear gain control (Heeger, 1992). It has since been proposed to be a canonical computation that is reused by multiple regions of the brain (Carandini & Heeger, 2012). Importantly, gain control serves an ecologically-motivated functional role in efficiently coding natural signals that have correlated statistical structure (Schwartz & Simoncelli, 2001)

A (divisive) normalisation term can be added to the above LN model, as well as several other parameters to specify the complete responses of the model neurons:

$$\lambda_d(s) = \gamma \frac{f_d(s)^\eta}{\sigma^\eta + f_\phi(s)^\eta} + \beta. \tag{4.3}$$

Here, an individual neuron's response is driven by the stimulus through $f_d(s)$, and this excitatory response is normalised by the terms in the denominator. One of those terms is the suppressive field, $f_\phi(s)$, which pools the responses of nearby neurons, making each neuron's response dependent on the activity of its neighbours. Because neighbouring neurons have stimulus-specific receptive fields, pooling the responses of neurons $f_\phi(s)$ is also a tuned function of the stimulus (by the same trick of treating synaptic inputs as a tuning function: $f_\phi(s) \approx \sum_j w_j s_j$). The suppressive field is typically more broadly tuned to stimulus features than the excitatory stimulus drive (Klímová et al., 2021).

Neurons also exhibit a typical sigmoidal pattern of response saturation (see Figure 4.1b for example) due to the divisive gain control. The responsiveness of the neural population is controlled both by σ, which shifts the range of stimulus strengths that are excitatory, and by γ which controls the overall response gain (or maximum spiking rate). Assuming that gain is evenly distributed among neurons, the population gain is expressed as $D\gamma$. Finally, the parameter β corresponds to a baseline level of neural activity which is not a function of the stimulus, and therefore reflects spontaneous background activity. This point will become relevant later, when examining how alpha affects stimulus-related processing, so it is worth noting now that Equation 4.3 implies an additive composition of stimulus-related and non stimulus-related (β) neural activity.

Representation in Spiking Population Codes

The equations so far specify the expected activity of neurons in the model population, in response to stimulus features like contrast and orientation. The spiking activity of neurons is then modelled as a homogenous Poisson process, using $\lambda_d(s)$ to determine the rate at which spikes are generated by each neuron[3]. Thus, neural activity exhibits apparent variability, or noise, while information is contained in the instantaneous rate of firing. To formalise this idea, take x as the vector of spike counts for a population of D neurons. For each neuron d, the likelihood of observing n spikes in a time window T is:

$$p(x_d = n|s) = \frac{(\lambda_d(s)T)^n e^{-\lambda_d(s)T}}{n!}. \tag{4.4}$$

[3]An assumption of the homogenous Poisson process is that the spikes of each neuron are independent, i.e. exhibit no covariance. The Poisson process also assumes spiking is 'memoryless' and captured only by an instantaneous rate. In reality, both of these assumptions are inconsistent with realistic neural spiking (which is both correlated and shows adaptation effects), but modelling these details, here, would yield little extra information for the added complexity.

The resulting model takes in a given stimulus as input and produces a sample of spiking activity from the population of neurons. This process approximates how visual neurons represent the sensory stimulus within a noisy evoked response that is distributed amongst the neural population. If the goal is to infer the stimulus from neural activity, then the collective firing of visual neurons can be interpreted as a probabilistic sample drawn from a distribution that encodes the likelihood of each possible stimulus. In principle, the likelihood distribution could be recovered by downstream neurons that perform a simple weighted summation of the visual population response[4] (Jazayeri & Movshon, 2006). This suggests that the brain could optimally decode the stimulus value and use the population activity to support a range of perceptual judgements that depend on probabilistic computations, like detection, discrimination, and optimal feature integration (e.g. combining noisy visual and auditory signals).

It is worth noting that early visual cortex contains many neurons that are arranged retinotopically and are tuned to a variety of visual features beyond orientation, like spatial frequency or binocular disparity (Hubel & Wiesel, 1968). Thus, a complete population response would certainly contain additional information that the brain could use to process rich naturalistic signals. In the present application, however, I focus on an abstract population that encodes orientation at a fixed location directly overlapping the stimuli presented in Chapters 2 and 3. Importantly, only orientation and contrast are varied in my experiments, so this hypothetical population contains all the behaviourally relevant information required to perform these tasks.

4.2 Perceptual Hypotheses and Actions

How might the brain use the information contained in the spike counts x, to guide action and perceptual decision-making about the unknown state, s, of the external world? This question can be approached from an abstract computational standpoint, rather than focusing on the specific circuits within the brain that perform decision-making and action. Ultimately, the goal is to understand how pulsed-inhibition might limit the available information within a biologically plausible sensory code, in a way that affects behaviour. Therefore, the visual encoding model derived above captures a sufficient level of biological detail, and realistic

[4] Recall that weighted summation, or linear filtering, is a basic computation of neurons, making this an attractive idea for how probabilistic decoding could be achieved in realistic neural networks.

decision-making processes that happen in the real brain can be modelled by a higher-level account (see Gold and Shadlen, 2007 for review of neural implementations of decision-making processes). Bayesian Decision Theory (BDT) provides a useful normative framework for understanding ideal decision-making in the face of uncertainty.

Perceptual decisions can be framed as an action selection problem in which the observer must choose an action, a, from a set of possible actions, A. Each possible action provides a certain amount of utility for the observer, which is a function of the state of the real world: $U(s, a)$, e.g. accurate actions tend to be most useful. BDT suggests that when the state of the world is unknown, the optimal action, a^*, is the one that provides the greatest expected utility:

$$a^* = \underset{a \in A}{\operatorname{argmax}} \, \mathbb{E}[U|x] = \underset{a \in A}{\operatorname{argmax}} \int U(s, a)\, p(s|x) ds. \tag{4.5}$$

Selecting the ideal action depends on two key factors: U, the utility of actions, and $p(s|x)$, the probabilistic information that spikes convey about the stimulus, i.e. the *posterior* distribution of s.

To apply the neural observer model to both the contrast detection and orientation estimation tasks, the utility function and posterior will need to be considered for each. Note that the posterior distribution is a function of x, which is the output of the encoding model. This suggests that the encoding model need not change between tasks because the visual system remains fixed, although its output will be variable on any given trial. In contrast, the utility function depends on the set of possible actions and will change from task to task. In the following sections, I will consider how to decode the posterior and how the actions and utility functions differ by task.

Bayesian Decoding of Spikes

The posterior distribution over the stimulus, s, given an observed sample of population spiking activity, x, can be calculated based on the modelled probability of spike generation, applying Bayes' theorem:

$$p(s|x) = \frac{p(x|s)\, p(s)}{\int p(x|s)\, p(s) ds} \tag{4.6}$$

In this expression, $p(x|s)$ is the likelihood of jointly observing all the spikes in x, i.e. the product of Equation 4.4 over all D neurons, $p(x|s) = \prod_d p(x_d|s)$. The prior

probability over stimulus is denoted by $p(s)$, and the term in the denominator expresses the marginal probability of x.

The activity of the modelled visual neurons is driven by both stimulus contrast, s_c, and orientation, s_θ, which makes the input to the encoding model multidimensional: $s = (s_c, s_\theta)$. Note also that realistic population codes also embed joint information beyond these two features. Despite this, behavioural experiments might probe only one of these feature dimensions, in isolation of the other. Perceptual hypotheses can still be informed using the posterior for jointly encoded stimuli, $p(s|x) = p(s_c, s_\theta|x)$. The posterior distribution about either orientation or contrast, alone, is obtained by marginalising the joint posterior over the dimension that is irrelevant to the task:

$$p(s_c|x) = \int p(s_c, s_\theta|x) ds_\theta, \qquad p(s_\theta|x) = \int p(s_c, s_\theta|x) ds_c. \tag{4.7}$$

This formalises how the same visual responses can be used to inform completely different perceptual-decisions by downstream brain regions, and that the structure of the task is what largely shapes behavioural responses. Although marginalisation here is performed using probabilities in a Bayesian context, it has been suggested that a realistic implementation in neural networks could be based on divisive normalisation circuits (Beck et al., 2011). This implies that focusing on orientation and contrast is akin to marginalising across other visual features (e.g. spatial frequency, binocular disparity, & wavelength).

Action and Utility in a Detection Task

In a detection task there are two possible world states, $s \in \{s_+ = \text{present}, s_0 = \text{absent}\}$. If the observer must respond both to presentations of present and absent stimuli (i.e. it is a forced choice task as in Chapter 2, and not a Go/No-Go design), then there are two available response options: $a \in \{\text{present, absent}\}$.

A deterministic utility function can be used that simply returns the accuracy of response choices against the true state of the stimulus:

$$U(s, a) = \begin{cases} 1, & \text{if } a \text{ is correct for } s, \\ 0, & \text{otherwise.} \end{cases} \tag{4.8}$$

Generally, utility could take on as many values as the number of contingencies that exist. Note that this particular function yields an equal utility for both correct rejections and hits, and this reflects the fact that an unbiased observer

would seek to maximise accuracy in both situations. Other utility functions are possible, for instance, if the task included a reward structure that incentivised either conservative or liberal detection responses more strongly.

Inserting this utility function into Equation 4.5 simplifies to an optimal decision rule that satisfies the need to maximise accuracy in the presence of uncertainty (see Ma et al., 2023, for its derivation):

$$a(x) = \begin{cases} \text{present}, & \text{if } p(s_+|x) > p(s_0|x), \\ \text{absent}, & \text{otherwise.} \end{cases} \tag{4.9}$$

Equivalently, this yields a simple rule in which the ideal observer would respond present if the posterior probability of the stimulus is greater than chance, i.e. $p(s_+|x) > .5$. Because the posterior probabilities depend on the prior probability of each stimulus class, the decision rule internally adjusts to situations where the prior odds of presenting a stimulus are unbalanced. That is, if there are more stimuli present than absent (or vice versa), then the decision rule will maximise accuracy if this is accounted for in the stimulus prior used in decoding.

Action and Utility During Orientation Estimation

In an estimation task, the observer must report a continuous stimulus rather a discrete choice. Accuracy will therefore depend on the discrepancy between the reported value a and the true stimulus value s. A utility function that rewards accurate estimation will seek to minimise this discrepancy and can be expressed in a general form:

$$U(s, a) = -|a - s|^{\rho}. \tag{4.10}$$

The exponent ρ controls how sharply utility decreases with estimation error. Larger exponents make utility decrease more quickly for larger errors, putting greater weight on avoiding large estimation errors.

Assume that the ideal observer always makes their response, $a = \hat{s}$, using the optimal action[5] in Equation 4.5. Then the Bayes estimator, $\hat{s}$, can be defined as

$$\hat{s} := \underset{\hat{s}}{\operatorname{argmax}} \int -|\hat{s} - s|^{\rho} p(s|x) ds. \tag{4.11}$$

Again, the Bayes estimator depends on the loss-exponent. For example, $\rho = 2$ produces a squared-error loss and suggests that the observer should report a as the

[5] A violation of this assumption would be that there is additional motor-related noise that alters the final report, which is something I explicitly modelled in Chapter 3 and found no evidence for.

posterior mean to achieve the best expected utility, while $\rho = 1$ yields an absolute-error loss leading to the posterior median as the value of a.

4.3 Pulsed-Inhibition of Neural Activity

The model so far describes how visual stimuli are encoded and decoded. To model oscillatory influences on behaviour, a final ingredient is needed to describe how alpha activity modulates neural processing. The working hypothesis of this chapter is that alpha oscillations represent pulsed-inhibition of ongoing neural activity (Mathewson et al., 2011; Mazaheri & Jensen, 2010; Schalk, 2015).

The pulsed-inhibition account suggests that alpha impacts visual processing by inhibiting sensory spiking activity. Direct evidence for this comes from observations that alpha power and phase reduce sensory spiking in macaque V1 (Dougherty et al., 2017) and sensorimotor cortex (Haegens et al., 2011). Evidence that alpha reduces spiking activity has also been seen in human intracranial EEG recordings (Chapeton et al., 2019). Beyond this, there is a large amount of behavioural and non-invasive neuroimaging evidence that is consistent with alpha oscillations inhibiting neural processing through a direct reduction in spikes (much of this was reviewed in the thesis' General Introduction).

While alpha oscillations can be characterised by separate oscillatory measures (like amplitude, phase, and frequency), the pulsed-inhibition account implies there is a singular inhibitory mechanism underlying alpha's effects on neural activity. This suggests that oscillatory measures could be used to recover the latent influence on neural spiking, by modelling the relative contributions of amplitude, phase, and frequency. From a modelling perspective, this is advantageous because it implies that oscillatory effects might be approximated by examining mostly the direct influence of inhibition on neural processing, while implicitly assuming the rhythmic and time-varying dynamics. This makes it important to understand the shape, or profile, of pulsed-inhibition's relationship to measures like EEG alpha amplitude and phase.

An important characteristic of pulsed-inhibition is that it has a non-zero mean, or baseline (see Figure 4.1c). This is captured in the idea that inhibitory suppression is delivered in pulses, instead of there being oscillations around (i.e., above and below) a base level of excitability. The alpha-band in EEG activity exhibits such an amplitude asymmetry and this could be explained by bouts of GABAergic feedback that silence pyramidal cell activity (Mazaheri & Jensen, 2010). Broadly,

this is consistent with other findings that instantaneous voltage and wideband local field potentials are better predictors of neural activation than measures of alpha-band power and phase that assume a zero-mean sinusoidal profile (Davis et al., 2022; Schalk et al., 2017).

To formalise pulsed-inhibition with a non-zero mean, we can model inhibition as a scalar function of time, t:

$$I(t) = A\left[1 + \cos(2\pi f t - \phi)\right]. \tag{4.12}$$

Here, A and ϕ are the amplitude and phase, respectively, of the pulsed-inhibition, and they may also be time-varying parameters that are governed by other neural factors that control alpha's generation and dynamics. This formulation encapsulates many of the core properties needed to model pulsed-inhibition, as inhibition of this form will oscillate around a non-zero mean, and return to uninhibited values when amplitude is low.

Next, consider how to recover pulsed-inhibition from the EEG signals observed on a given trial. If it is assumed that inhibition oscillates in the alpha band (i.e. $f \approx 10$Hz is fixed) and that the alpha amplitude, A_t, and phase, ϕ_t are measured at time t, then inhibition can be inferred as a latent state. This is achieved by rearranging Equation 4.12 to be a function of amplitude and phase, rather than time:

$$I(A_t, \phi_t) = gA_t\left[1 + \cos(\phi_t - \phi^*)\right]. \tag{4.13}$$

Here, ϕ^* is introduced to denote the phase at which inhibition peaks. This compensates for the fact that measured EEG alpha phase could differ from the phase of latent pulsed-inhibition because the interpretation of EEG signals is reference-dependent. Moreover, such differences can arise because EEG sums the activity of many electrical sources in the brain that have differing orientations due to the folding of the cortex (Nunez & Srinivasan, 2006). As a result, EEG-measured phase can differ from phase measured intracranially at the source. In Chapter 2, the inhibitory phase ϕ^* was able to be fixed based on the inhibitory relationship between oscillatory phase and subsequent visually evoked potentials.

There are some indications that EEG alpha waveforms directly reflect pulsed-inhibition at the neural level, such as the fact the alpha-band has a non-zero mean as predicted by putative pulsed-inhibition (Mazaheri & Jensen, 2010). However, it is not appropriate to directly place EEG values into the model, because the natural units of EEG (μV) differ from the units that define spiking rates (e.g. spikes per second, or Hz). Thus, the free parameter, g, is introduced in Equation 4.13 to

control the unknown scaling between alpha amplitude and the strength of latent pulsed-inhibition.

In the following sections, I have chosen to model the effect of inhibition using the amplitude of alpha oscillations and ignoring their phase. This simplification is justified by the pulsed-inhibition model, which assumes that amplitude and phase measures have a unitary underlying effect via inhibition. In some applications of the model, it is also only feasible to fit and compare performance in two contrasting alpha power bins, instead of the continuous modulation needed to account for alpha phase. Ignoring phase information comes at the cost of reducing the temporal specificity of modelled perceptual effects. However, it is clear from the pulsed-inhibition model that the effects of amplitude are the average of the effects of phase over the oscillatory cycle. Moreover, estimating perceptual effects using only amplitude will result in a conservative estimate, as amplitude ignores the stronger peaks of inhibition that can be tracked when accounting for phase.

The pulsed-inhibition model I have formulated here is a special case of the more general equation that was explored in the behavioural modelling used in Chapter 2. The general equation allows for sustained, or weakly phase-modulated inhibition, and the data from Chapter 2 support that inhibition is only partially released at higher alpha amplitude. The details of this finding, while important for characterising the profile of inhibition, do not conflict with the current choice to focus solely on amplitude. Rather, they suggest that amplitude-only modelling can account for a large proportion of alpha's inhibitory effects.

4.4 Modelling the Neural Contrast Response Function

A detection task is a natural way to explore the dynamics of the neural model in the context of pulsed alpha-band inhibition. This is because the modelled neural responses exhibit well-established non-linear dynamics that are driven by stimulus intensity (Albrecht & Hamilton, 1982) and which are captured by divisive normalisation (Equation 4.16; Heeger, 1992). The contrast response function maps out these dynamics as a function of stimulus contrast. The pulsed-inhibition model makes a concurrent claim that neural responses are affected also by internal time-varying inhibitory states.

In Chapter 2, I explored this joint interplay between neural sensitivity in response to stimuli and alpha power and phase. There I showed that pulsed-inhibition reduces the baseline and response gain of hypothetical neural responses,

using a phenomenological model of behaviour, i.e. the psychometric function. Note that the modelling in Chapter 2 uses signal detection theory (SDT) to link psychometric function changes to a latent decision variable, under the assumption that it is corrupted by Gaussian noise. In contrast, the visual encoding model defined above offers a biologically inspired neural framework that serves as a more plausible substrate for decision-relevant information, while adopting realistic assumptions that noise originates from neurons' Poisson variability and the structure of the population. The relevant question now is whether these behavioural dynamics from Chapter 2 are similarly captured within the biologically-inspired neural model that I have outlined here.

In Silico Experiment and Model Predictions

It is possible to investigate the effects of inhibition on detection performance using a purely simulation-based approach. To achieve this, the computational nature of the model is exploited by running simulations where pulsed-inhibition is targeted at various aspects (or parameters) of the model. Detection behaviour is then read-out from the population code, as explained above, as a function of stimulus contrast intensity. Precise simulations can be readily performed to investigate the influences of neural inhibition upon putative behavioural outcomes, whereas analogous electrophysiological experiments may be resource-intensive or difficult to implement.

The resulting analyses are a valuable tool and provide insight into the model's behaviour prior to fitting or comparing any models quantitatively. The simulation experiment is therefore a *prior predictive* model comparison, as opposed to a *posterior predictive* comparison which would be based on indices of fitted models, conditional upon observed data (Gabry et al., 2019). Given, the computational complexity involved in fitting the neural model to real data, the *prior predictive* analysis can be used to rule out implausible models, saving resources from being wasted in ill-posed or *a priori* uninformative models. It is tempting to dismiss these predictive simulations as a precursory step before model fitting; however, they give core insights into the theory that the model expresses, and the outputs constitute major findings that inform the theoretical link between alpha oscillations and perception. For instance, as the forthcoming results will show, some possible pulsed-inhibition effects are, in fact, facilitative for behaviour, suggesting they are implausible at face value as the relationship between alpha and detection behaviour is known to be inhibitory (e.g., (Mathewson et al., 2009, 2011)).

Table 4.1 outlines the possible targets of pulsed-inhibition, denoted here as I, within the model. While Chapter 2 explores similar inhibition targets within a Signal Detection Theoretic model, it is interesting to note that some of the inhibition models here overlap with the SDT model, while others are entirely unable to be captured by the basic SDT model. This supports the idea that the mechanistic model reveals insights that are not apparent otherwise. Specifically, orientation tuning and the suppressive field gain of the model are not captured when using SDT, because of its simplistic basis in an abstract decision signal with Gaussian noise and threshold relation to behaviour.

Inhibition target	**Relevant parameter**	**Inhibition model**
Response gain	γ	$\lambda_d(s) = I\gamma\frac{f_d(s)^\eta}{\sigma^\eta+f_\phi(s)^\eta} + \beta$
Baseline firing	β	$\lambda_d(s) = \gamma\frac{f_d(s)^\eta}{\sigma^\eta+f_\phi(s)^\eta} + I\beta$
Contrast sensitivity/gain	σ	$\lambda_d(s) = \gamma\frac{f_d(s)^\eta}{(I\sigma)^\eta+f_\phi(s)^\eta} + \beta$
Response exponent	η	$\lambda_d(s) = \gamma\frac{f_d(s)^{I\eta}}{\sigma^{I\eta}+f_\phi(s)^{I\eta}} + \beta$
Orientation tuning	κ	$f_d(s) = e^{I\kappa\cos(s-s_d^*)-I\kappa}$
Suppressive field gain	$f_\phi(s)$	$\lambda_d(s) = \gamma\frac{f_d(s)^\eta}{\sigma^\eta+If_\phi(s)^\eta} + \beta$

Table 4.1: Possible targets of inhibition within the neural observer model

Inhibition, denoted I, is modelled as a multiplicative scaling or adjustment of the parameters of the model.

The neural observer model was first fit to human response data. Figure 4.2b shows the aggregated contrast response function from the detection experiment in Chapter 2 and that it is well-captured by the decoded detection responses from the fitted neural model. As a second step, the model is inhibited in line with the possible inhibition models outlined in Table 4.1, shown in Figure 4.2c–h. In comparing the simulated contrast response functions, and the various effects of pulsed-inhibition, the goal is to understand which possible inhibitory targets capture the pattern of human performance expressed in Figure 4.2a under low and high alpha power conditions.

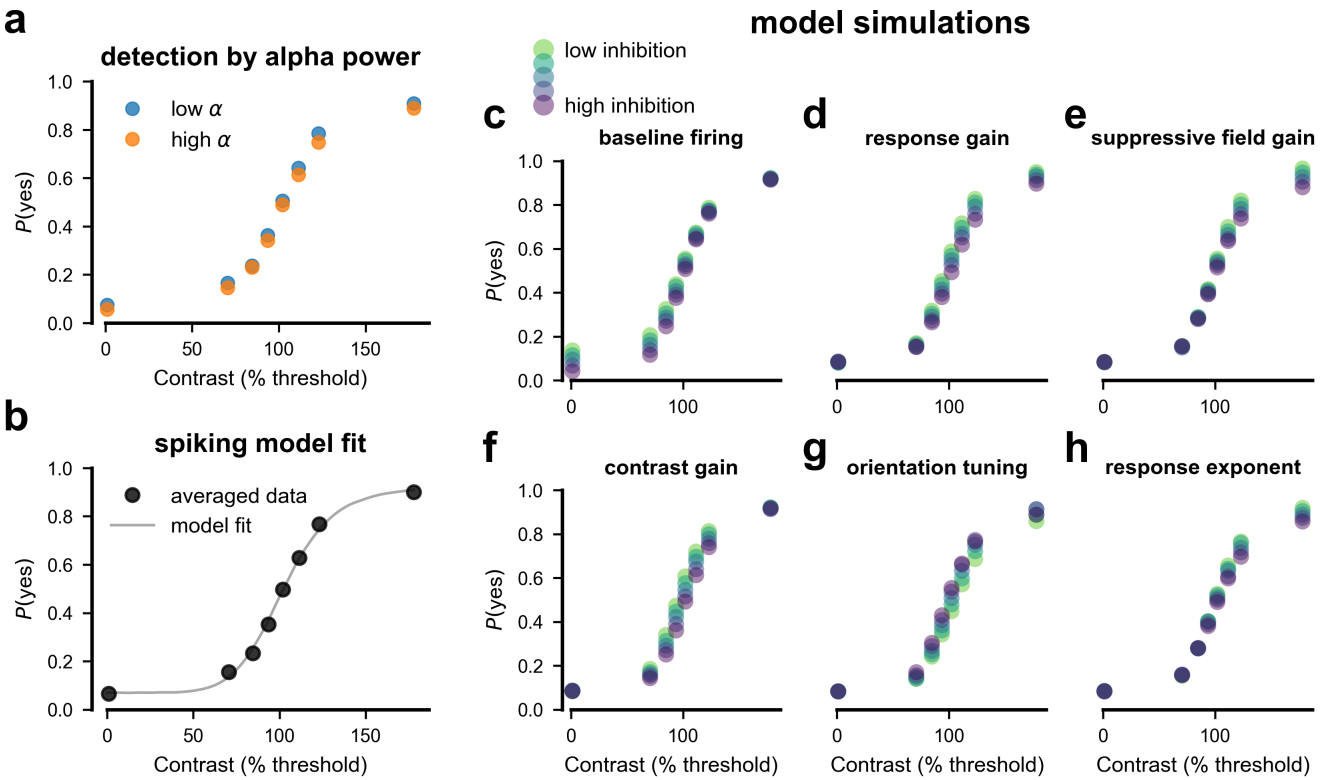

Figure 4.2: Simulated inhibition of a spiking model's contrast response function

a Human response data from the detection experiment reported in Chapter 2, split by low and high prestimulus alpha power (the bottom and top quartiles are shown in blue and orange, respectively). The ordinate shows the proportion of 'yes' responses. Contrast is shown on the abscissa in percentage units of each observers' individual threshold, to allow aggregation of the contrast response functions. **b** The neural observer model was first fitted to the overall detection data, aggregating alpha power conditions, to find plausible parameters for simulation analyses. **c**–**h** Predictive simulations showing the characteristic impacts of inhibiting six different parameters of the neural observer model. Darker colours indicate stronger inhibition, which corresponds to the neural state during higher alpha power or stronger pulsed-inhibition. The strength of simulated inhibition was arbitrarily chosen to roughly equate the performance modulation across simulations.

Baseline Activity Inhibition Model

Real data from human observers show a decreased probability of positive detection responses across the contrast response function when alpha power is high. Notably, no single simulated inhibition model can capture this full pattern of results. The baseline inhibition model (Figure 4.2c) is the only inhibition target that captures the detection decrement at low, sub-threshold, contrast levels. Detection responses to zero-contrast stimuli correspond to the false alarm rate. This is a key quantity used in estimating SDT sensitivity and criterion metrics.

Experimental findings show that alpha power modulates primarily decision confidence and SDT criterion, rather than perceptual accuracy (i.e. SDT sensitivity; Iemi et al., 2017; Limbach and Corballis, 2016; Samaha et al., 2017b). Because of this, previous theoretic work based in SDT has posited that the mechanistic action of alpha oscillations is to reduce sensory noise that can bias decisions (Samaha et al., 2020). Intuitively, detection tasks may be biased towards more frequent 'yes' responses when alpha power is low because of increased spontaneous perceptual noise that is confused for task-relevant evidence by a decision-maker. In line with this theory, inhibiting the baseline parameter in the model suppresses the spontaneous component of neural responses driven by noise.

A recent study investigated the influence of alpha power on the psychometric function for detection (Pilipenko & Samaha, 2024), and showed a consistent downward shift in performance across the psychometric function. The results of this study are similar to those of my detection task in Chapter 2, which are shown in Figure 4.2a. Pilipenko and Samaha (2024) argued that these findings are captured by the baseline excitability model, in which alpha only modulates spontaneous activity. However, the simulation results presented in Figure 4.2c speak against this and it is clear that inhibiting spontaneous noise alone is insufficient to capture the performance modulation at higher contrast values. For stimulus-related signals to inform perceptual decisions, they must eventually dominate the noise levels of the visual system as contrast-intensity increases. Therefore, at high contrast values the noise must be comparatively weak relative to the stimulus-related activity. As such, the effect of noise on detection becomes negligible at higher contrast, consistent with the pattern observed in the simulation (Figure 4.2c).

To my knowledge, these are the first results that present a neurally specific and explicit analysis of neural excitability effects on perception. Previous work has largely been in the context of abstracted SDT analyses used on detection tasks that only presented stimuli at a single (threshold) intensity. Thus, the model simulation presented here is a valuable tool for investigating the neural plausibility of the theorised baseline excitability effects. Previous work has modelled baseline modulation in the psychometric function using accuracy (Pilipenko & Samaha, 2024), and not a latent perceptual space like that of SDT or the spiking activity of my neural observer model. It is important to remember that accuracy is an observable high-level summary of behaviour that is influenced by neural processing across many distinct stages and regions. As a result, there may be subtle influences

of pulsed-inhibition on the latent processes that produce behaviour which are not easily discovered when analysing observed patterns of raw behaviour. This is especially relevant when investigating the contrast response function because of its nonlinear effects on internal representations.

Response Gain Inhibition and Suppressive Field Models

What captures the detection decrement across the entire range of psychometric functions, if not the baseline inhibition model alone? The simulation reveals three models that produce suppression of above-threshold performance: response gain, suppressive field gain, and response exponent models. It is likely that pulsed-inhibition modulates a combination of baseline activity and one of these other models.

Response gain in the neural observer model corresponds to any factors that multiplicatively modulate neural firing rate and effectively limit the total output capacity of individual neurons in response to fixed input (Ferguson & Cardin, 2020; Silver, 2010). The effects of response gain inhibition are shown in Figure 4.2d which exhibits the characteristic modulation of performance at the upper asymptote of the contrast response function. Modulation is greatest in the above-threshold stimulus region because facilitation (or, opposingly, depression) of neural responses multiplicatively interacts with stronger stimulus intensities, which already boost neural firing rates. Likewise, response gain changes are weakest below threshold because the stimulus-driven component of neural activity is small or non-existent. Given the explicit connection between neural firing and the pulsed-inhibition model, response gain is a natural and plausible account of the effects of oscillations on detection performance. Furthermore, a previous study has shown evidence for a response gain model accounting for alpha power's effects on detection (Chaumon & Busch, 2014).

A similar effect on the contrast response function arises in the suppressive field model simulations, shown in Figure 4.2e. In this model, inhibition boosts the contribution of suppressive fields in the denominator of the normalisation equation (i.e. $f_\phi(s)$ in Equation 4.3). The formulation of this model is, to my knowledge, a completely novel application. Suppressive fields have primarily been modelled in the spatial and pattern vision literature, where a key question is the influence of competing stimuli on neural responses and psychophysical detection (Carandini, 2004).

Suppressive fields explain a range of visual response properties, like cross-orientation or surround suppression (Carandini & Heeger, 2012; Xing & Heeger, 2001), as arising from lateral inhibition by nearby neurons with receptive fields for competing stimuli[6]. It is therefore unusual to model the contribution of suppressive fields to the contrast response function of a single stimulus. Despite this, a recent study has shown that prestimulus alpha power promotes the perception of the tilt illusion (Williams et al., 2024), which arises from lateral inhibition in early visual cortex. This implies that alpha-band inhibition may operate through lateral inhibitory connections, a notion further supported by neurophysiological findings that indicate that attentional rhythms arise from inhibitory competition between receptive fields (Kienitz et al., 2018). These findings justify further examination of pulsed inhibition mediated by the model's suppressive fields, even though, in the detection task from Chapter 2, the suppressive fields did not receive any exogenous input.

Increasing the drive to suppressive fields has a similar effect on the model's contrast response function to that of decreasing the response gain. This can be shown mathematically using Equation 4.3, in which the component governing neural response drive, $f_d(s)$, is divided by a term that includes the suppressive field, $f_\phi(s)$. Assuming that contrast values are substantially above threshold, the suppressive field overpowers other terms in the denominator allowing σ to be ignored, i.e. $f_\phi(s) \gg \sigma$ (because $f_\phi(s)$ depends on contrast). The response of model neurons under response gain inhibition is approximately:

$$\lambda_d \approx \frac{I f_d(s)}{f_\phi(s)}, \tag{4.14}$$

where I is an inhibitory scaling factor that decreases the neural response drive if $I < 1$. This is equivalent to inversely scaling the term in the denominator:

$$\frac{I f_d(s)}{f_\phi(s)} = \frac{f_d(s)}{I^{-1} f_\phi(s)}, \tag{4.15}$$

as is the case if suppressive field drive is increased. Interestingly, the dependence of suppressive field drive on higher stimulus contrasts can be seen in the simulations. Comparing the upper asymptotes of Figure 4.2d with Figure 4.2e, it can be seen that response gain modulation decreases at the highest contrast, while suppressive

[6]Suppressive fields refer to a general concept in computational visual neuroscience (which is also captured by the divisive normalisation computation; c.f. Carandini, 2004), but the implementation is not achieved solely by lateral, or horizontal, connections (Angelucci et al., 2017).

field modulation starts at higher contrast levels and shows no apparent decrease as contrast moves away from threshold values. This suggests that the two mechanisms are dissociable, in principle, by their dependence on σ.

Finally, the model in which inhibition reduces the response exponent, η, effectively reduces the slope of the contrast response function (Figure 4.2h). The exponent controls the sensitivity of neural responses to increasing stimulus intensities and has a profile similar to that of the suppressive field inhibition model. This could be due to the fact that all terms in the divisive normalisation in Equation 4.3 are raised to an identical exponent, which preferentially increases the influence of the larger terms. Consequently, the effect of modulating the exponent may, to a large extent, replicate the suppressive field effects, since pooling over larger receptive fields makes a substantial contribution to the normalisation equation. Consistent with this, experimentally induced lateral inhibition in mouse V1 produced slope changes in neural contrast sensitivity (Del Rosario et al., 2025). Furthermore, it is worth noting that the exponent captures a largely phenomenological, as opposed to mechanistic, aspect of the model. Some approaches assign different exponents to each term in the normalisation to achieve a better fit, but this comes at the risk of making the model overly flexible. Therefore, specifying a single exponent is a necessary simplification that supports parsimony. For these reasons, the exponent modulation is excluded from further analyses in favour of the suppressive field model, which produces similar yet more interpretable changes in slope.

Contrast Gain and Tuning Width Models

The contrast gain inhibition model (Figure 4.2f) shows the expected modulation which is a right-ward shift of the contrast response function under inhibition. This horizontal translation is evident only around threshold-level contrast values, in the portion of the response function that is approximately linear. The modulation can be likened to increasing the threshold value of the contrast response function. Contrast gain changes are readily observed in neural recordings as an effect of contrast adaptation (Heeger, 1992; Sclar et al., 1989). Such changes effectively shift the sensitive range that neurons are responsive to[7] as though the effective contrast of the stimulus has been reduced (Silver, 2010). Model comparisons have previously favoured response gain over contrast gain (Chaumon & Busch,

[7]The dynamic, or sensitive, range of a neuron is the portion of its input-output response function where the response is approximately linear.

2014) and there is little to suggest that contrast gain changes could account for the empirical pattern of detection responses at face value (Figure 4.2a). This is because the performance modulation is restricted to regions around the threshold-level of performance, where the contrast response function is linear. Therefore, the contrast gain inhibition model can be excluded from further analysis of the neural observer model.

The effect of simulating tuning width changes due to inhibition is shown in Figure 4.2g. Remarkably, the effect here appears to be like that of the response gain model but in the opposite direction: tuning width inhibition improves contrast detection at higher contrasts. In creating the simulation, I have assumed that the effect of inhibition would be to broaden the width of individual neurons' tuning curves, resulting in worsened neural selectivity. This is broadly consistent with observations that coding precision is typically better as a consequence of increasing attentional gain (Martinez-Trujillo & Treue, 2004; Reynolds & Heeger, 2009) and worsened under conditions of energetic stress, such as food scarcity (Hung et al., 2025; Padamsey et al., 2022). Changes in the tuning width of individual neurons are also not a natural target for the model because orientation selectivity is largely contrast invariant (Albrecht & Hamilton, 1982; Sclar & Freeman, 1982) due to gain control from divisive normalisation (Carandini & Heeger, 2012).

Why would broader orientation tuning facilitate detection responses in the model? Orientation selectivity is effectively a constraint on the range of stimulus values that elicit a response from an individual neuron. If that neuron's becomes less selective, then it is more likely to exhibit responses to nearby orientations, and, likewise, nearby neurons are more able to contribute spiking responses to stimuli further from their own preferred orientations. The net effect is that there are more neurons available to respond to a given stimulus. This will produce more orientation noise in the population, making it harder to precisely decode orientation but easier to decode whether a stimulus was present. Thus, improved detection arises at a similar tradeoff seen in how higher baseline activity, which is spontaneous noise, also facilitates detection responses at the cost of higher false alarms.

Model Fitting

It is clear from the previous section that the neural observer model could incorporate pulsed-inhibition in a variety of ways. The simulation results suggest many of these possible models can be ruled out *a priori* except for two models: a *baseline*

plus response gain inhibition model and a *baseline plus suppressive field* model. In theory, both models might account equally well for human response data and simulations alone will not yield insight into which model is more suitable than the other. Further comparison of these two models requires fitting them to the actual response data, so that the quantitative evidence for each model can be compared.

The two fitted models are plotted against the detection responses in low and high alpha power conditions of Chapter 2 in Figure 4.3a–b. Both models were freely fitted containing identical numbers of parameters, with all parameters allowed to vary. As expected, both models capture the decrease in detection performance across the entire range of the psychometric function. Bayesian Information Criterion (BIC) scores were used to compare the quantitative fit of the models. The baseline and response gain inhibition model had a BIC of 128.42 ±0.55 standard deviations (sd). The baseline and suppressive field model had a BIC of 128.85 ±0.74 sd. A model without inhibitory dynamics was fitted, serving as a null model, which had a higher BIC score of 140.88 ±1.12.

BIC values approximate the marginal evidence for a model, allowing them to be converted into an approximate Bayes Factor (Wagenmakers, 2007), which compares models based on the ratio of marginal likelihoods. Bayes Factor analyses revealed that both models were substantially more likely than the null (BF_{10}'s > 408). However, there was only negligible to weak evidence favouring the baseline and response gain model over the baseline and suppressive field model (BF = 1.24). The model comparisons are summarised in Figure 4.3c using model probability weights (computed from BIC; cf. Wagenmakers and Farrell, 2004) showing that the two models are roughly equiprobable.

The equal fit of the two models suggests that pulsed-inhibition may be just as well explained by response gain inhibition as it is by suppressive field changes, at least as far as the available data are able to reveal. One possible explanation is that strength of baseline inhibition trades off against the characteristic shape of modulation that each model can explain at higher contrasts. That is, because the inhibition of response gain is more apparent around threshold-level contrast (see again Figure 4.2d) compared with suppressive field modulation, the fitted baseline inhibition could be reduced to achieve a comparable fit to the data. Indeed, this is what is seen in Figure 4.3d, which plots the parameter estimates for both models. The baseline parameter, β, is inhibited more strongly in the model with suppressive modulation, which makes up for the weaker inhibitory modulation at

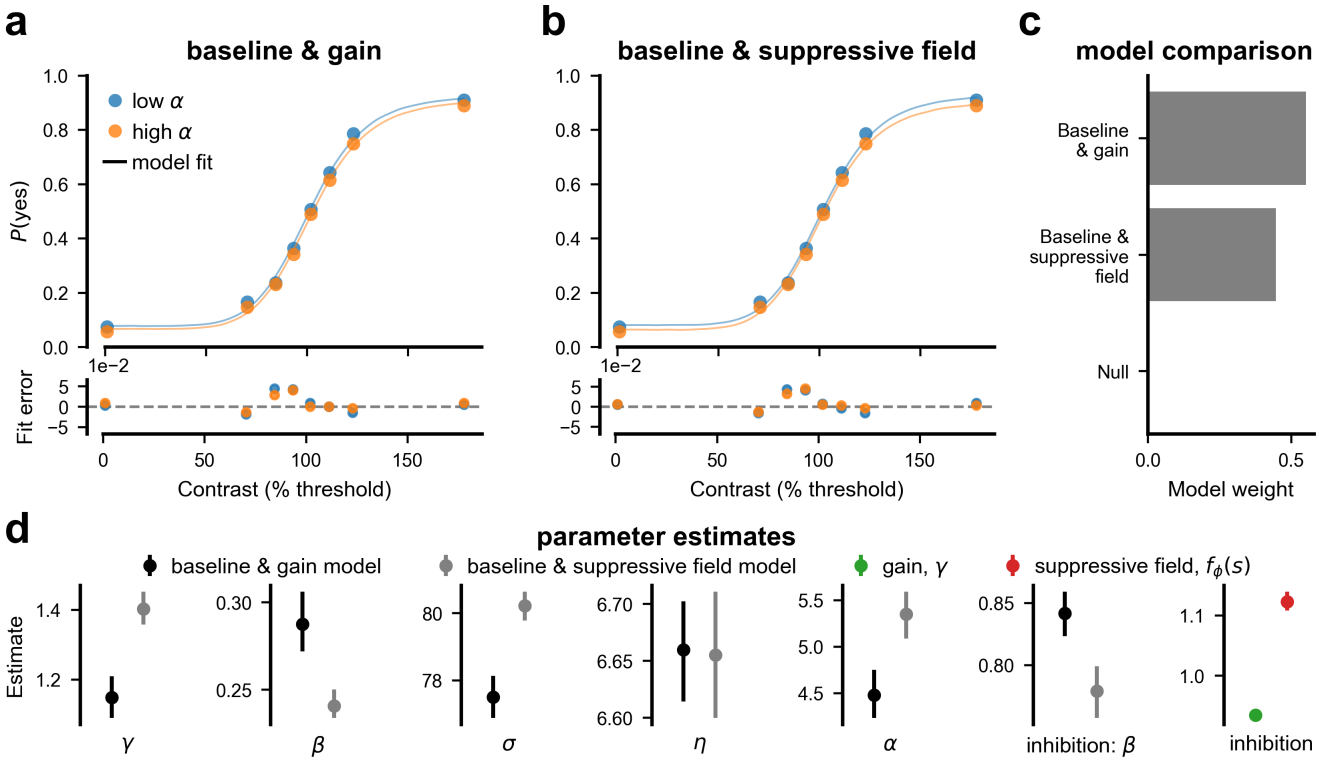

Figure 4.3: Fitted models of inhibited contrast responses

a–b Contrast response functions for two fitted models in which pulsed-inhibition modulates the neural observer model. Continuous lines show the response functions predicted by the models, which were generated by Bayesian decoding of 300,000 simulated neural population spike counts at each of 30 contrast values across the function. The bottom panel shows the error between model predictions at the contrast values with observed data points. **c** Model probability weights (calculated from Bayesian Information Criteria), including a null model that was fit to the same high and low alpha power detection data, but without inhibition. Higher model weights indicate a more probable model, under the assumption that all models are equally plausible. **d** Parameter estimates for the two models shown in **a–b**. A bootstrapping procedure was used to calculate standard deviations of the estimates (vertical lines). The parameter symbols correspond to the model shown in Table 4.1 except for α (which was an additional free parameter that is unrelated to pulsed-inhibition and was used during fitting to control the gain of pooled suppressive fields, $\alpha f_\phi(s)$, relative to normalised input drive). The final two rightward-most panels show inhibition parameters (normalised to 1 = no inhibitory effect; values < 1 show inhibition shrinks the targeted parameter value).

threshold contrast levels[8].

In an attempt to mitigate over-flexibility in the model fitting, models were fitted using an alternative approach. This time, the main parameters governing the contrast response function were fitted first to the low alpha power data, and fixed at those values. Then, in a second step, this model was inhibited in line with both competing models, to account for the high alpha power condition. The result is that the uninhibited contrast response function remained fixed across models, which is consistent with the assumption that pulsed-inhibition affects only the high power condition. This resulted in negligible to weak evidence for the baseline and suppressive field model over the baseline and response gain model (BF = 1.47, model probability = .59)[9].

In summary, here, a neural observer model was fitted to human detection behaviour. A simulation-based approach was used to explore how inhibiting the model affects its contrast-driven responses across the psychometric function. A clear result is that inhibition must modulate baseline firing rate in combination with either neural response gain or suppressive field strength. While quantitative model fitting was unable to neatly adjudicate between these two possible models, the results are nonetheless informative and suggest that the models have a highly similar impact on the behavioural result. The possibility that alpha oscillations strengthen the contribution of suppressive fields in divisive normalisation has not previously been considered, yet it is consistent with evidence the alpha oscillations act through lateral connections involved in surround suppression (Harvey et al., 2013; Williams et al., 2024). To resolve this, future experiments could examine whether alpha modulates contrast-related perceptual processing when multiple stimuli compete in the population representation (Busse et al., 2009). Such experiments would be expected to show modulation in line with the suppressive field account but not if the response gain account is correct.

4.5 Modelling Efficient Neural Codes for Orientation

In Chapter 3, I showed that alpha oscillations influence orientation estimation and promote repulsion away from the cardinal orientations that dominate natural

[8]Note that many other parameters also change between models.

[9]An alternative fitting method could be considered, in which the baseline inhibition parameter is fixed before fitting the free parameters for response gain vs suppressive field, forcing the baseline inhibition to be constant across models. However, it does not resolve the problem, in principle, because it is still unknown whether the fixed baseline inhibition is correctly 'unmixed' from the true effects of response gain or suppressive field modulation.

signals. To apply the neural observer model to these results, it is necessary to first understand more about efficient population codes. For the present purposes, efficient coding provides a principled way to adapt the spiking population model so that it can represent the non-uniform natural statistics of orientation.

To begin, the efficient coding hypothesis proposes that neural systems maximise the shared information between neural activity and environmental signals (Barlow, 1961). Shared information can be quantified in information theoretic terms as mutual information, but a more common application to neural coding is the Fisher Information (Wei & Stocker, 2016). Fisher Information reflects how sensitive a neural code is to local variations around a given stimulus. It is a key quantity in theoretical neuroscience because it provides a common unit of analysis that bridges across representational systems, regardless of their form (e.g. representations in behaviour, probabilistic codes, or neural geometry).

Fisher Information is a function of the stimulus, $J(s)$, and can be expressed in a general probabilistic form, as well as in terms of the neural observer model,

$$J(s) = \mathbb{E}\left[\left(\frac{\partial}{\partial s}\log p(x|s)\right)^2 \middle| s\right] = \sum_d \frac{\lambda_d'(s)^2}{\lambda_d(s)}. \tag{4.16}$$

The righthand side of the equation can be obtained by substituting $p(x|s)$ using Equation 4.4. This shows that for a population of model neurons with independent Poisson noise, the Fisher Information depends on each neuron's firing rate and its derivative.

By default, the Fisher Information of the neural observer model is uniform over orientation, $J(s) \approx \text{const}$, reflecting the fact that its representational coverage is uniform. This is because the neurons in the population have identically shaped tuning curves and are shifted copies, tiled evenly across orientation space. However, natural environments have a prior distribution over orientation that reflects the greater prevalence of vertical and horizontal orientations in natural scenes (Girshick et al., 2011). The goal under efficient coding is to match the representational capacity of the neural code to the prior distribution of orientation. This would result in a neural code whose (square-root) Fisher Information is proportional to the prior (Wei & Stocker, 2015),

$$\sqrt{J(s)} \propto p(s). \tag{4.17}$$

There are many ways that a population code could achieve this by altering any combination of the individual neurons' tuning preferences, tuning widths,

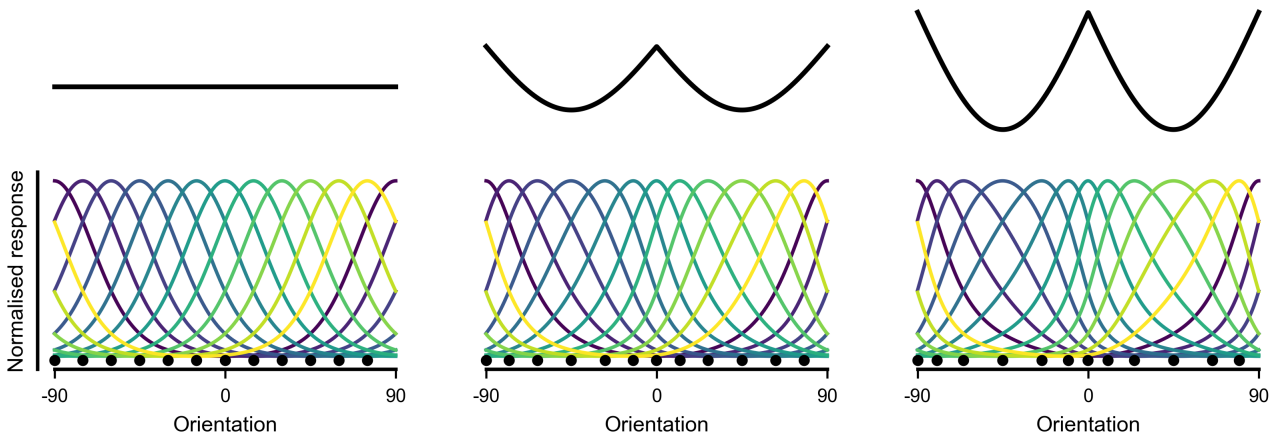

Figure 4.4: Warping of tuning curves under efficient coding of the stimulus prior

Three examples are shown of population codes that efficiently represent (circular) stimulus spaces with different prior distributions. The prior distributions are shown by the solid black line above each population, and increase from uniform (left) to increasingly bimodal (right). Individual neurons are differently coloured with black dots showing their most-preferred stimulus.

or response gains. An elegant and theoretically-motivated solution is to warp the stimulus space in proportion to the prior distribution (Ganguli & Simoncelli, 2014). The outcome of this is that the density of neural tuning, in terms of tuning preferences and shapes, is naturally allocated by the prior, while retaining the tiling property of the population (see Figure 4.4). To achieve this, consider a transfer function, $\Gamma(s)$, that maps the stimulus space to efficient sensory space where the prior distribution is now uniform, $S \rightarrow \tilde{S}$. The sensory space transformation is

$$\tilde{s} = \Gamma(s), \qquad \Gamma(s) = \int_{-\infty}^{s} p(u)\, du, \tag{4.18}$$

which implies, in simple terms, that the transfer function of orientation into sensory space is warped according to the cumulative distribution of the prior. Notice that the derivative of this function is equal to the stimulus prior, $\Gamma'(s) = p(s)$. This ensures that the transformed sensory space will expand around areas where the stimulus is frequent and shrink where it is less frequent.

Efficient neural tuning functions $f_d^*(s)$ can be constructed by feeding the transformed stimulus $\tilde{s}$ into the existing model tuning functions $f_d(s)$ that assume a uniform distribution of inputs,

$$f_d^*(s) = f_d(\Gamma(s)). \tag{4.19}$$

The result is that the neural model is now reparameterised by a function that embeds the prior distribution into the tuning curves. By the change of variables, the Fisher Information of the efficient code is the original uniform $J(s)$ scaled by the derivative of the transfer function,

$$J^*(s) = J(\Gamma(s))[\Gamma'(s)]^2, \tag{4.20}$$

which satisfies the efficient coding constraint from Wei and Stocker (2015)

$$\sqrt{J^*(s)} \propto p(s). \tag{4.21}$$

Simulated Orientation Biases Under Pulsed-Inhibition

Recall that in Chapter 3, I showed that perceived orientations were repulsively biased away from the cardinal axes (horizontal/vertical), and that this repulsive bias increased under conditions of high alpha power. Using the neural observer modelling framework developed here, I ran simulations to examine how different forms of inhibition could give rise to orientation estimates that are likewise more strongly repulsively biased. These simulations represent a *prior predictive* analysis, consistent with the strategy I used above to explore the model's predictions under different mechanistic targets of pulsed inhibition. To achieve this, the neural observer model was constrained by warping the tuning functions based on a generic orientation prior of the form $p(s) \propto 2 - |\sin(\theta)|$ (Girshick et al., 2011; Wei & Stocker, 2015). The resulting efficient neural population was then inhibited in four ways: (1) reduced baseline activity, (2) reduced gain, (3) stronger normalisation by suppressive fields, and (4) increased tuning width.

The effects of each simulated inhibition type are shown by plotting the resulting tuning curve for a single neuron in Figure 4.5a and the derivative of that neuron's tuning is plotted below in 4.5b. Note that the Fisher Information of the neural code is computed using the former divided by the latter (see Equation 4.16). Computing the total Fisher Information as a function of hypothetical alpha power reveals that Fisher Information increases under baseline inhibition and decreases under all other types of inhibition (Figure 4.5c). In Chapter 3, I found that the Fisher Information computed from observers' behaviour is instead decreased by alpha power. This suggests that baseline inhibition, by itself, cannot account for the alpha-effects seen in real human orientation judgements from Chapter 3. Intuitively, Fisher Information increases when baseline activity is inhibited because this activity is not orientation-specific and, therefore, adds orientation

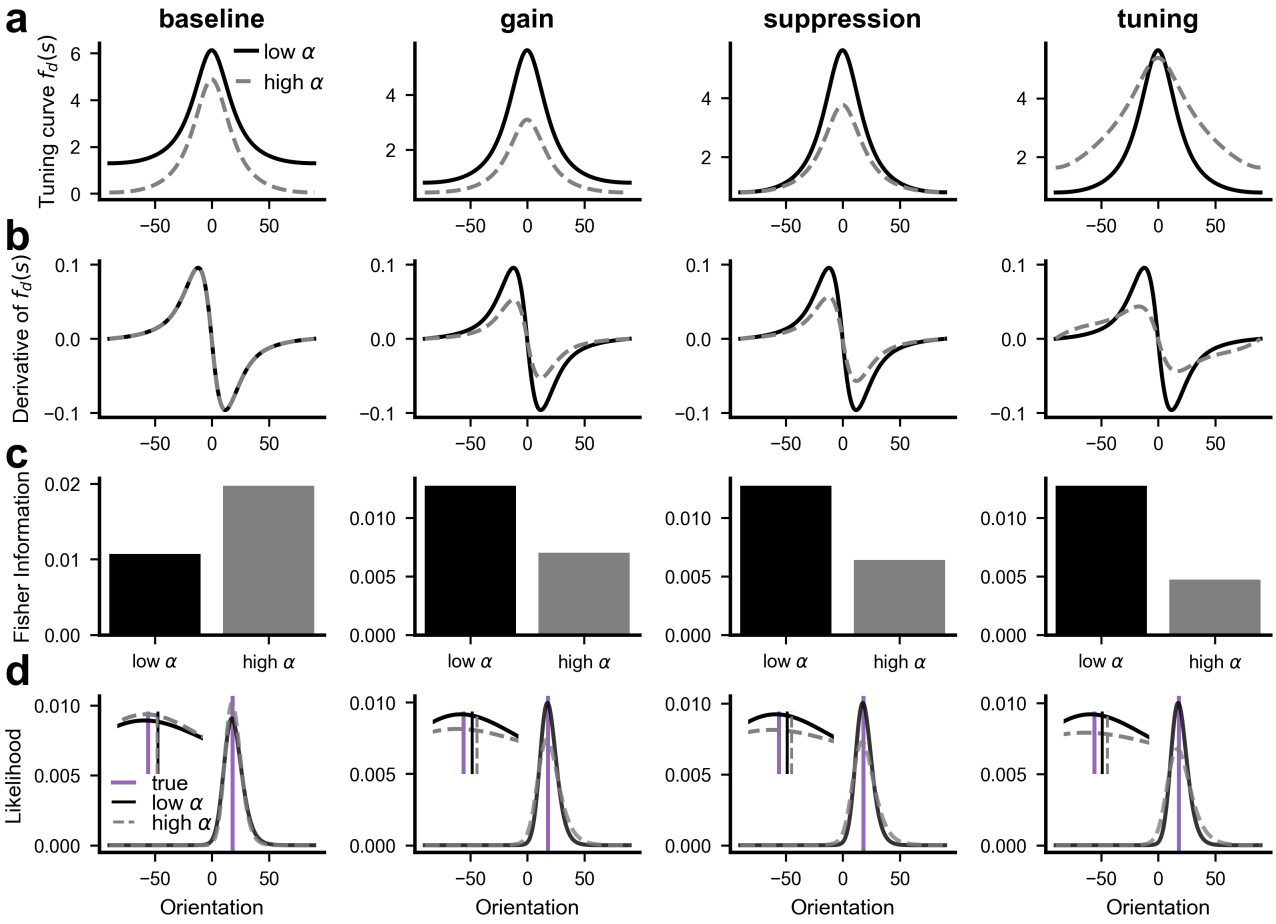

Figure 4.5: Orientation biases as a function of inhibiting baseline activity, gain, suppression, and tuning

a Tuning curves of a single neuron in the model population, with a preferred stimulus at 0. **b** The derivatives of the tuning curve in **a**. When calculating Fisher Information, the squared derivative is divided by the value of the tuning curve. **c** The Fisher Information of all neurons in the population, averaged over all orientations. **d** A single stimulus (purple line) was presented to the neural populations and the likelihood of the stimulus was decoded from the resulting spiking activity, and averaged over 10,000 presentations. The inset shows a zoomed view of the distributions near the stimulus. The solid black and dashed grey lines show the decoded posterior mean for low and high alpha power, respectively. Decoded orientation that fall to the right of the purple line are estimates of the stimulus that are repulsed away from the high prior density at 0.

noise to the population response. This can also be seen through the fact that denominator in Equation 4.16 shrinks with the reduced firing rate, whereas the derivative is unaffected by baseline inhibition (Figure 4.5b).

To investigate orientation biases, each population code was used to produce 10,000 simulated responses to each single orientation. The neural likelihoods were computed and averaged over the simulations (Figure 4.5d). Examining

these likelihoods reveals that all the inhibition models produce an estimate of the stimulus that is repulsively biased away from the zero orientation. This is because the models produce neural likelihoods that are skewed away from the high prior density of the stimulus at 0 and 90°. However, the baseline inhibition model does not produce stronger repulsion during high alpha power. Notably, the gain, suppression, and tuning models all demonstrate increased repulsion in high alpha. A computational explanation is that all of these inhibition models produce greater sensory noise when inhibited, which amplifies the skew in the neural likelihood distribution.

It is clear from this simulation analysis that it is difficult to select a single model as a sole explanation for alpha's effects on orientation biases. Nonetheless, it does clarify that the mechanism responsible for producing stronger repulsive biases is an increase in neural likelihood noise (as in Wei and Stocker, 2015), which would be reflected in reduced Fisher Information over all orientations[10]. Any inhibitory effect that reduces orientation-related activity will produce stronger repulsive biases. Inhibiting response gain and the suppressive field can achieve this by reducing the total amount of spikes that are emitted, making it harder to decode orientation precisely. In contrast, increasing the tuning width does not reduce population spiking (Figure 4.5a), suggesting that this mechanism can be ruled out because it would not reduce total firing rates as has been shown for alpha oscillations (Dougherty et al., 2017; Haegens et al., 2011).

Fitting the Observer Model to Behaviour

To fit the neural observer model to behaviour, it is possible to exploit the fact that Fisher information expresses the representational capacity of both human observers and the model's orientation code. Therefore, the model was fit to the Fisher Information estimated from behavioural responses in Chapter 3, rather than being fit directly to raw behaviour. The shape of the transfer function $\Gamma(s)$ was allowed to vary in the fitting procedure, as well as unknown population gain and inhibition strength. Inhibition was modelled as a simultaneous reduction in response gain and the baseline activity of the population. Although Figure 4.5 shows that baseline inhibition produces cardinal attraction rather than repulsion,

[10] It is interesting to note that noise in the neural likelihood need not correspond to noise in the population spiking activity. Inhibiting baseline firing removes spontaneous noise activations from the population response, but this does not affect the width of the neural likelihood distribution. To change likelihood noise (i.e. width/uncertainty), neural activity that is *informative* about orientation must be altered.

it was included for consistency with the contrast response function modelling that supports its involvement in pulsed-inhibition.

Note that, in this setting, suppressive field modulation yields effects that are effectively the same as response gain (up to a multiplicative constant). This is because, with only a single stimulus, the two mechanisms can only be distinguished by their nonlinear behaviour across different contrast levels. Consequently, the orientation estimation task—carried out at just one contrast level—does not provide enough information to meaningfully discriminate between the suppressive field and response gain accounts. Therefore, the response gain was chosen because it is the most plausible mechanism to date given its support in the past literature (Chaumon & Busch, 2014).

Figure 4.6a shows the Fisher Information from behaviour in low and high alpha power, and the fitted model. The strength of response gain inhibition is fitted by the model and is the only parameter that controls the difference in Fisher Information between alpha power conditions. The fitted population code was presented with 20,000 stimuli and the presented orientation was decoded from each generated sample of spiking activity. The estimation errors produced by the model are plotted in Figure 4.6c and show the expected pattern of stronger repulsion from cardinals during the simulated high alpha power condition. The results of this fitting and simulation procedure confirm that reduced spiking activity in an efficient visual code is a sufficient explanation for the stronger orientation biases seen during high alpha power.

Summary

The neural observer model was adapted to produce repulsive orientation estimation biases, which are a characteristic of efficient representational codes. Simulating the effects of inhibition revealed that baseline inhibition cannot produce stronger biases, like those seen in human behaviour. This contrasts with the apparent necessity of baseline inhibition to explain contrast detection results. The reason for this divergence is that the two tasks are informed by separate information, and the baseline activity does not convey consistent orientation signals but can falsely signal stimulus presence. It is likely that response gain or suppressive fields may explain the effects seen in human behaviour, but it is not possible to dissociate these models based on the present estimation data. Putting aside the specific mechanisms that are targeted by inhibition, the modelling confirms

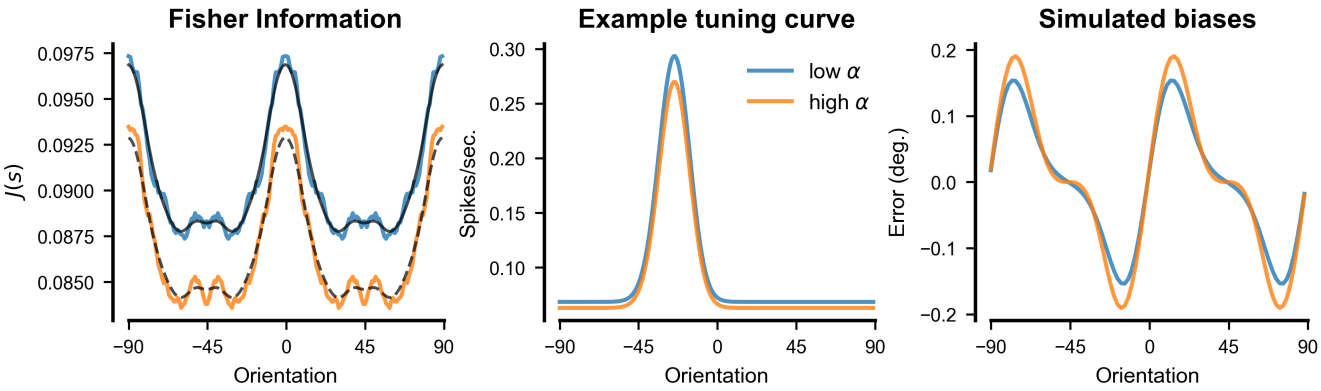

Figure 4.6: Repulsive biases in an orientation code fitted to behavioural Fisher Information

a The Fisher Information from a behavioural orientation estimation task as a function of high and low prestimulus alpha power. The behavioural data were averaged across vertical and horizontal, exploiting the symmetry of orientation space. Dark lines show the Fisher Information of the fitted neural code, only neural inhibition accounts for the difference between high and low alpha power. **b** An example tuning curve from the fitted population neural code. Inhibition was modelled as a reduction in both the response gain and baseline activity of the population. **c** Spiking activity was simulated from the model and decoded using an ideal neural observer model that performs orientation estimation. Estimation error (reported minus true orientation) is plotted showing a stronger repulsive bias away from the peaks of the Fisher Information at 0 and 90°, which are prevalent orientations in natural environments.

that reduced spike rates are an essential component that explains how perceptual performance is more strongly biased during high alpha power.

4.6 Discussion and Implications of the Model

Creating a neural observer model of rhythmic visual perception is a way to formalise the theory of alpha's impact upon perception. The model is a compact, and necessarily quantitative, tool that can be applied to existing results, as I have done for the detection and orientation judgements presented in Chapter 2 and 3, respectively. Importantly, the structure of the model and its core assumptions about pulsed-inhibition do not change to accommodate the differences in these perceptual tasks. Strong theories should be generalisable and accommodate a range of findings with minimal adjustment to their core assumptions. The key premise of the current model is that alpha oscillations reflect pulsed-inhibition and the effect of this is to reduce the strength of encoded visual information.

The perceptual effects of inhibited visual information are not immediately obvious. Functional properties of visual neurons influence the content of encoded

information and its relevance to distinct perceptual tasks. For instance, normalisation and gain control make the contrast response function nonlinear, and this affects how pulsed-inhibition modulates detection at different levels of contrast. This is clear when considering how the above analyses showed pulsed-inhibition must target at least two components that control visual representation strength (baseline firing and either response gain or suppressive field gain). Of the various modulation effects that were simulated, no single parameter could account for the suppression of detection response across the full psychometric function over contrast. This fact would not be obvious when reasoning about empirical results without recourse to the latent neural processes that determine perception, which are formalised in the model. Overall, the model suggests that pulsed-inhibition may target baseline neural responses in addition to either neural response gain or suppressive normalisation.

The model reveals a further subtlety in that the effects of baseline inhibition, while required to explain alpha's effects on detection, are not able to produce stronger repulsive cardinal biases in orientation perception. The findings above show that baseline inhibition should, instead, be facilitative and reduce biases. This is due to the reduction in spontaneous noise activity. Specifically, in the context of an orientation estimation task, spontaneous activity contributes noise information that corrupts the accuracy of the perceptual decision (i.e. orientation judgement). Interestingly, this is true also in detection, as spontaneous activity corrupts the perceptual decision (i.e. stimulus detection) and provides more frequent false alarms, but also an increased hit rate. Thus, inhibiting baseline activity (i.e. the spontaneous or non-stimulus driven component of visual activity) is facilitative in orientation perception but removes accurate responses in detection tasks. This underscores the difficulty of synthesising the results of different perceptual tasks and shows that model predictions can be nuanced even when effects can be attributed to a common latent source.

Across the two tasks that were examined, the model accounts for human behavioural reports on the basis that pulsed-inhibition reduces the spiking rates of visual neurons that encode stimulus properties. The target of spike-rate inhibition is both the stimulus-driven activity and the baseline responses, suggesting that the effect of pulsed-inhibition is applied to all components that contribute to the population response; whether exogenously driven or spontaneous. This is compatible with neurophysiological evidence that V1 spiking activity is coupled to alpha phase when spiking is stimulus-driven and also during periods without stimulation (Dougherty et al., 2017). Analysing the neural observer model revealed,

first, the possibility that stimulus-driven activity modulation reflects a pulsed-inhibition of neural response gain. Importantly, the model also suggests a novel second hypothesis that pulsed-inhibition suppresses stimulus-driven activity via divisive normalisation computations.

Divisive normalisation inhibits neural responses in a stimulus-tuned manner, which operates through the suppressive fields that pool stimulus-driven activity of neighbouring neurons in the population. These suppressive fields are implemented by long-range inhibitory connections within the cortex and recurrent inter-regional feedback (Angelucci et al., 2017). Consequently, the suppressive effects of pulsed-inhibition may also be stimulus-tuned. Thus, a major prediction of the model is that pulsed-inhibition should be relevant to the processing of joint stimuli, such as multiple oriented gratings presented at once. Potentially, alpha oscillations are supported by the circuits that mediate inhibitory interactions between the neural responses of multiple stimuli. This hypothesis is important because there are many neuronal and perceptual findings that arise from such circuits, and these have been well-characterised in past visual neuroscience studies (reviews in Carandini and Heeger, 2012 and Carandini et al., 2005). Therefore, it may be possible to extend a theory of pulsed-inhibition's effects on perception, based on incorporating existing theories of visual functions that are not based in oscillations.

For example, surround modulation is well-studied in visual neuroscience (Angelucci & Bressloff, 2006; Angelucci et al., 2017) and is a phenomenon in which perceptual and neural responses are inhibited by stimulation outside the excitatory receptive fields. This is a key neuronal mechanism that could provide fruitful advances in characterising pulsed-inhibition and the alpha rhythm in perception. For example, there are several characteristics of surround modulation that provide testable hypotheses. First, surround modulation is tuned to features like orientation and spatial frequency, suggesting that pulsed-inhibition's effects should be prominent for these same features and similarly tuned. Second, human psychophysical data have shown that surround modulation increased with eccentricity and was not present for foveally presented stimuli (Petrov et al., 2005). This suggests an intriguing hypothesis that the effects of pulsed-inhibition may not be observed for stimuli presented in central vision. Variation of oscillatory perception across the visual field has not been considered and could account for differences in studies that find alpha phase effects with peripheral stimulation (Busch et al., 2009; Harris et al., 2018; Zazio et al., 2021) and those that do not find them

with central stimuli (e.g. Ruzzoli et al., 2019). Nonetheless, the suppressive field account suggests a principled reason to expect visual field variation.

In addition to surround modulation, another inhibitory phenomenon is cross-orientation suppression which is also based in divisive normalisation. Cross-orientation arises when a mask grating with varying orientation and contrast is superimposed over a target grating, suppressing responses in receptive fields even though the contrast and orientation of the target feature remain unchanged. Importantly, cross-orientation suppression presents spatially overlapping stimuli and the mechanisms that mediate this type of inhibition may arise from different circuitry to surround modulation. The former may arise from feedforward circuitry (such as synaptic outputs from LGN), and the latter from feedback and horizontal inhibitory connections in visual cortex (Angelucci et al., 2017). Such experiments could help in understanding whether the alpha oscillations that affect visual function are primarily a reflection of thalamic versus cortical sources, with the idea being that feedforward inhibition in cross-orientation interactions arises in efferent signals of the thalamus. In theory, this could yield a principled way to dissociate the network origins of pulsed-inhibition if alpha were shown to modulate the effects of one type of inhibitory circuit over the other.

Moving beyond whether pulsed-inhibition modulates response gain or suppressive fields, further predictions can be made by examining the impacts of expanding certain assumptions of the model. For instance, it was assumed that the neural population describes a fixed spatial location that overlaps the stimuli presented to probe perception experimentally. Future work could extend the model to include multiple populations that are distributed across retinotopic space. Thi would yield an image-computable neural observer model, which would make the suppressive fields an explicit function of spatial location and allow the above surround modulation hypotheses to be modelled. If pulsed-inhibition impacts the encoding strength of spatially distributed visual stimuli, then this model predicts that spatial discriminations such as Vernier acuity should also deteriorate under high alpha power (as borne out in data from our lab; Carman, Harrison, Bex, & Harris, in prep.). Interestingly, size-contrast illusions (such as the Ebbinghaus illusion) have been shown to correlate with resting state alpha power, such that higher power promotes stronger illusory size modulations reflecting greater influence from the surround spatial context (Chen et al., 2020). Extending the current modelling approach using the spatial dimension has some support already and could yield insights into how alpha oscillations impact visuo-spatial processing in a range of conditions, potentially including the processing of naturalistic images.

4.7 Summary

I have shown in this chapter that it is possible to construct a detailed computational model of visual sensory neurons that explains how encoded visual information is shaped by alpha oscillations. In the model, alpha oscillations reflect a rhythmic inhibitory effect that suppresses spiking activity. Using Bayesian Decision Theory, I investigated how alpha-based inhibition can shape the information conveyed in oscillating neural activity and thereby affect perception. In contrast detection, the inhibition of response gain and baseline firing yields modelled psychometric functions that match behaviour (as shown empirically in Chapter 2). Furthermore, efficient neural representations show biases when firing rates are inhibited and there are fewer spikes available to code visual orientation (as shown for behaviour in Chapter 3). Integrating ideas about the inhibitory effect of alpha oscillations could explain other results, like the tilt illusion's modulation with alpha (Williams et al., 2024), using a more extensive model of visual processing and also suggests model-based approaches to understanding broader oscillatory effects, such as in spatial vision. The modelling also suggests a novel hypothesis, namely, that inhibition might act through the suppressive fields used in divisive normalisation, but new experiments are needed to investigate this possibility.

Chapter 5

Rhythmic Visual Perception Arises from a Surround Modulation Mechanism

Perception is increasingly understood as rhythmic rather than continuous, reflecting cycles of excitation and inhibition in the brain (VanRullen, 2016). Neural oscillations are believed to play a computational role in the coordination of information processing within distributed assemblies of neurons, with rhythmic activity emerging from the circuit-level dynamics of networks that feature inhibition (Womelsdorf et al., 2014). Although previous work has emphasised how neural oscillations affect the sensitivity of the neural system to the strength of perceptual information (Mathewson et al., 2011), visual perception is also shaped by contextual computations that determine the content of what is perceived. Inhibitory neuronal mechanisms shape neural responses through surround modulation, allowing responses to a stimulus to be influenced by the spatial context in which it appears (Angelucci et al., 2017). In natural vision, spatially correlated inputs activate many neurons simultaneously, and surround modulation can sharpen neural representations by suppressing redundant or competing information (Carandini & Heeger, 2012). Currently, it is unknown whether rhythms in perceptual sensitivity emerge during the processing of spatial context and potentially shape the broader structure of visual experience.

Alpha-band (8–14 Hz) activity has been associated with reduced neural excitability and behavioural fluctuations in detection (Mathewson et al., 2011), sup-

This chapter is written for submission in a Brief Reports format.

porting the view that oscillations regulate the timing and routing of perceptual information (Jensen & Mazaheri, 2010). More recently, there has been growing evidence suggesting oscillatory modulation of centre-surround mechanisms in perception. For example, Williams et al. (2024) recently reported that the tilt illusion (in which the perceived orientation of a central stimulus is biased by an oriented surround stimulus) becomes stronger when stimuli are presented during periods of high-amplitude posterior alpha activity. This could indicate that alpha rhythms play a role in functional suppression in local centre-surround circuits. Consistent with this, stimulation of regions outside the population receptive fields of human V1 produced alpha power within the receptive field (measured with electrocorticography and functional magnetic resonance imaging; Harvey et al., 2013). Furthermore, recordings in awake macaques have demonstrated rhythmic (3–6 Hz) inhibition between competing stimuli that fall within each other's suppressive surround regions (Kienitz et al., 2018).

These findings suggest that rhythmic inhibitory dynamics and spatial surround computations may reflect shared cortical mechanisms. However, it is presently unclear whether rhythmic fluctuations simply gate visual excitability or arise within the feature-selective architecture of the visual code itself. If the latter is the case, this implies the radical notion that neural oscillations are part of a canonical computation that alters the geometry, and not simply gain, of how sensory codes respond to structured inputs.

5.1 Modelling and Empirical Findings

To illustrate these ideas, we built a simple network oscillator model in which oscillations arise in coupled excitatory and inhibitory neural populations (Figure 5.1a). Orientation tuning was modelled using excitatory subpopulations that respond to preferred orientations. Surround modulation was incorporated as long-range horizontal projections that synapse with local inhibitory subpopulations (c.f. Fitzpatrick, 2000; Figure 5.1b). The model captures cortical inhibition which is mediated through GABAergic interneurons that modulate neural gain (Ferguson & Cardin, 2020) and are also implicated in the generation of low-frequency cortical oscillations (Jones et al., 2000). Importantly, stimulating local inhibitory subpopulations resets the phase of ongoing rhythmic activity in excitatory populations (Figure 5.1c; Canavier, 2015). This models how a stimulus presented in one spatial location (e.g. an oriented surround) could shape the dynamics of distant

neural populations through inhibitory signals that spread laterally (Jones et al., 2000).

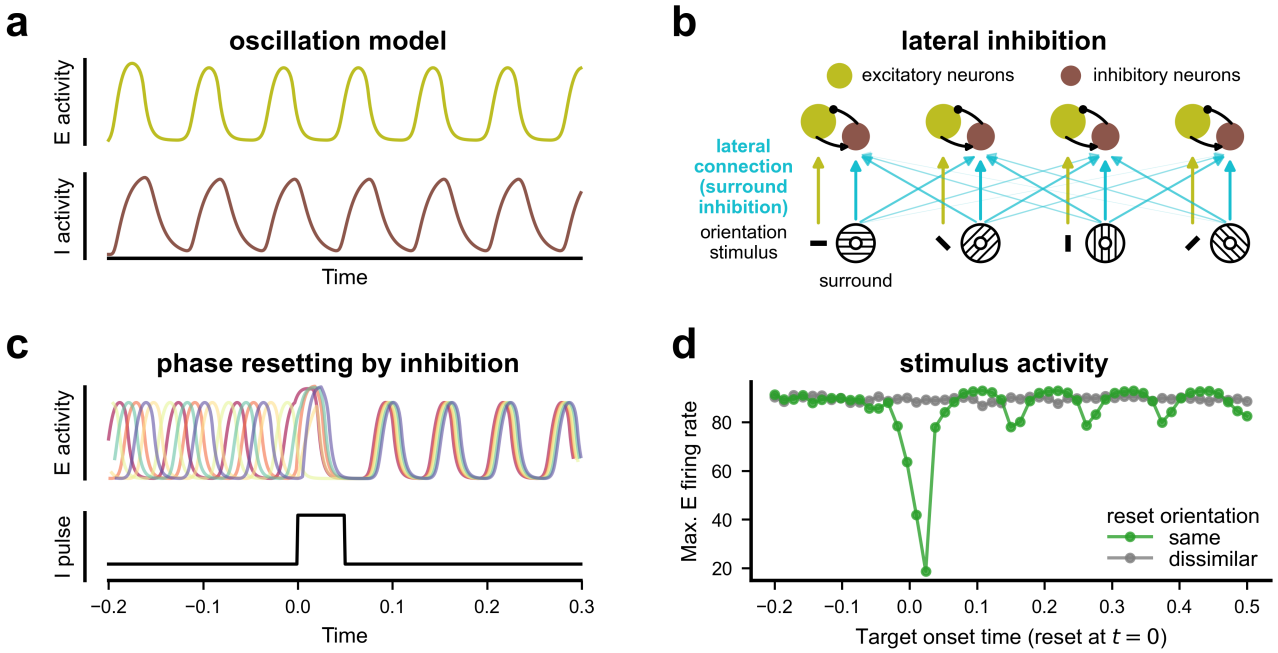

Figure 5.1: Network oscillator model with lateral inhibition

a Simple network oscillator model involving coupled excitatory (E) and inhibitory (I) neural populations (Wilson & Cowan, 1972). **b** Lateral inhibition network. Neural subpopulations respond preferentially to oriented stimuli, and long-range horizontal connections project to local inhibitory subpopulations that mediate surround inhibition. Lateral connections (blue lines) are strongest for similar orientations (shown by line width), producing orientation-tuned inhibition. **c** Phase resets are induced in excitatory activity (top panel) by stimulating inhibitory units with an impulse (bottom panel). **d** Simulated stimulus responses. An excitatory stimulus was simulated at various latencies relative to a simulated surround stimulus that resets oscillations at time zero. Coloured lines show the maximum stimulus-driven excitation (a proxy for stimulus detection) as a function of the target's orientation relative to the simulated surround orientation.

A hallmark of surround modulation is that it is orientation-tuned and strongest when the surround stimulus matches the orientation of the central stimulus (Angelucci et al., 2017). Because lateral connections in our model are similarly orientation-tuned, exogenously driven phase resets will be prominent among populations that are tuned to similar orientations and receive inhibition from the surround. Figure 5.1d shows that simulating oriented inputs at different latencies relative to an inhibitory surround input produces oscillations in the strength of stimulus-driven excitatory activity. This is strongest when the surround is similar in orientation to the preferred orientation of the excitatory subpopulations. The model therefore captures how surround modulation, which is mediated by local

inhibitory neurons, can produce orientation-dependent oscillations in the neural activity that encodes an external stimulus.

Having established the plausibility of these mechanisms, we conducted a behavioural experiment to examine the prediction that surround modulation produces rhythmic inhibition. Human observers performed a detection (i.e. 'present/absent') task where they responded to an oriented grating that was briefly presented on 50% of trials (Figure 5.2a). Importantly, the probe grating appeared at various onset asynchronies (-100 to 1000 milliseconds) relative to a separate surround grating. This allowed us to examine the latent time course, across trials, of observers' detection ability relative to the onset of the surround using a *dense sampling* paradigm that is commonly used to measure rhythms in behaviour (Kienitz et al., 2021). As expected, detection accuracy was diminished when the probe stimulus was presented close in time to the surround stimulus, due to visual masking (Figure 5.2b). This was followed by a period of improved detection that exhibited robust rhythmicity (Figure 5.2e).

We computed time-resolved Signal Detection Theory (SDT) sensitivity and criterion measures using a sliding window analysis (Figure 5.2c; see *Methods*). We then analysed the frequency content of the resulting time course data, excluding times prior to 100 ms (relative to surround onset) to avoid confounding any transient effects related to the initial visual masking. This revealed a ~5 Hz oscillation in perceptual sensitivity (Figure 5.2f) that appeared to explain the similar oscillatory peaks in the raw accuracy time series (Figure 5.2e). There was no apparent low-frequency oscillatory peak in the criterion time course, suggesting that the surround stimulus induced rhythms in perceptual sensitivity and not decision-related biases.

While it is interesting that behavioural rhythms arise using a visual surround as a 'reset' stimulus, the crucial prediction from our modelling is that perceptual oscillations are induced by orientation-tuned surround modulation. Our experiment was able to test this because the surround and central probe stimuli were randomly iso-oriented or cross-oriented (e.g. orthogonal) on each trial (see Figure 5.2a). In theory, any differences in perception between these orientation conditions must reflect surround modulation that is tuned to oriented features, as opposed to generic rhythmic temporal attention effects (Landau, 2018). Therefore, the time-series of behaviour in the cross-orientation condition, in which tuned inhibition would be weakest (e.g. Angelucci et al., 2017), is an important control condition against which performance in iso-oriented timeseries must be compared.

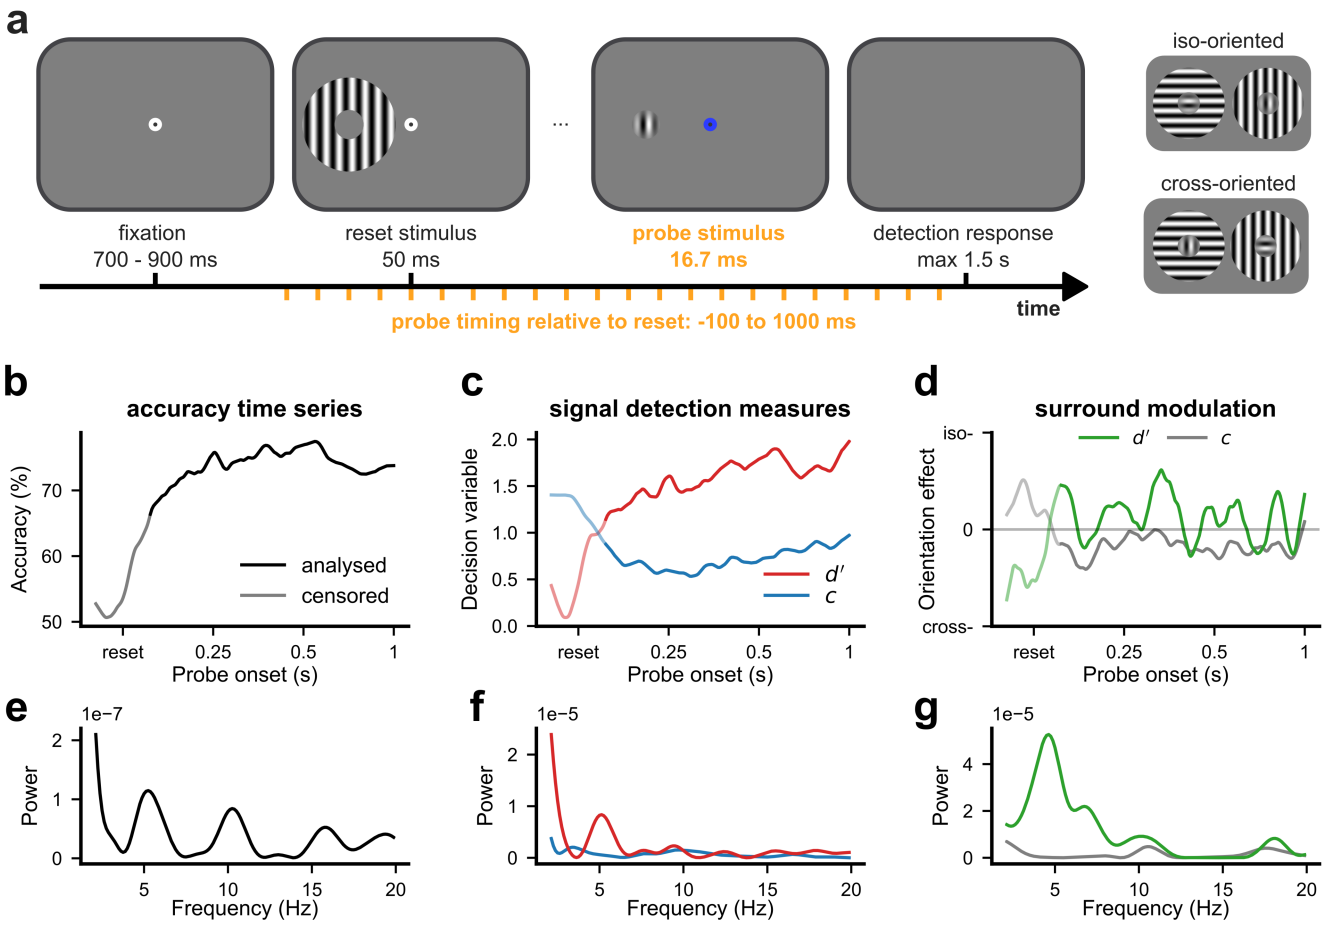

Figure 5.2: Orientation-dependent surround modulation oscillates at a theta frequency

a Trial sequence of the visual detection task. Peripheral probe target stimuli were presented at a latency relative to a spatial surround stimulus across trials. The orientations of the surround and probe stimuli were counterbalanced and could be iso-oriented or cross-oriented (orthogonal; right panels). **b** Timecourse of accurate responses (percentage hits and correct rejections) computed from a sliding window analysis. Time series data within ±100 ms of the surround reset stimulus was omitted from frequency analyses. **c** Signal Detection Theory sensitivity and criterion time series (light colours indicate omitted data as in **b**). **d** Time series of the orientation-tuned surround modulation effect, showing whether sensitivity or criterion were higher for iso- or cross-oriented stimuli. **e–g** Low-frequency oscillatory power of accuracy, signal detection, and surround modulation time courses, present in **b**, **c**, and **d**, respectively.

To address our main hypothesis about the orientation-tuning of rhythmic perception, we computed an index of the orientation-tuned effect based on the difference between iso- and cross-oriented time courses (shown in Figure 5.2d for both sensitivity and criterion). This revealed a prominent theta frequency oscillation in sensitivity (Figure 5.2g), suggesting that the orientation alignment

between the suppressive surround and the target underlies the theta rhythm in perceptual sensitivity. To quantify the reliability of this effect, we used a bootstrapping approach to compute resampled timeseries and their corresponding frequency spectra. We then fit each resampled spectrum using a spectral fitting algorithm (Donoghue et al., 2020) that allowed us to assess both the confidence in detecting oscillations and the amplitude and frequency variability inherent in the collected data. This method isolates oscillations and corrects for possible aperiodic structure that is a major concern in analysing behavioural rhythms (Brookshire, 2022; Harris & Beale, 2025). A low-frequency oscillation was detected on 98.4% of bootstrap samples with the most commonly detected frequency in the theta band (see Figure 5.3).

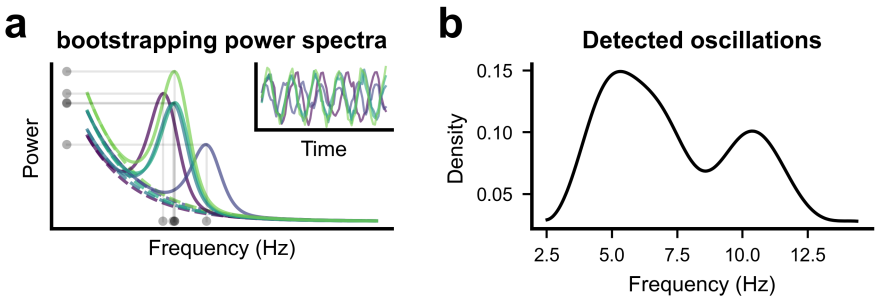

Figure 5.3: Bootstrap estimates of oscillatory surround modulation

Bootstrap resampling was used to estimate the variability of oscillations in the time course of the surround modulation sensitivity effect (see Figure 5.2d) **a** Illustration of bootstrapping process. Oscillations were detected by fitting the resampled power spectra, separating aperiodic activity and oscillatory peaks. Inset shows example time series data. **b** Kernel density estimation of the oscillatory peak frequency detected across bootstraps. Oscillatory peaks were detected in 98.4 % of bootstrap samples.

In summary, we used a psychophysical detection task to examine rhythms in perception that result from the surround modulation properties of neural coding. We show that the onset of an oriented surround stimulus evokes a low-frequency oscillation in perceptual contrast sensitivity, which exists at times extending beyond the transient effects of surround inhibition. This oscillation is specific to perceptual sensitivity, as opposed to decisional bias and, crucially, depends on the orientation of the surround stimulus. The comparison between iso-oriented and cross-oriented surround conditions is crucial to linking perceptual rhythms with surround modulation mechanisms, due to the well-characterised orientation-tuning that exists in surround modulation circuits (Angelucci et al., 2017).

We propose that rhythmic perception arises naturally from two aspects of inhibitory visual circuit processing that have, until now, been treated as dis-

tinct. In visual cortex, horizontal neuronal connections that drive inhibitory interneurons are necessary to explain the orientation-tuned surround suppression of receptive fields (Angelucci et al., 2017; Fitzpatrick, 2000). Concurrently, inhibitory interneurons are also responsible for the oscillations in cortical circuits that modulate excitatory gain control (c.f. Womelsdorf et al. (2014)). We modelled these inhibitory dynamics using simulated neural populations that express lateral-inhibition, allowing them to be orientation-tuned. The result is that orientation-tuned inhibitory populations can drive perceptual rhythms by resetting the phase of local oscillatory dynamics. The account that we formalise in our modelling provides a compelling mechanism by which neural oscillations in cortex could reflect the same circuits that implement lateral-inhibition and underlie surround modulation.

Our results suggest a mechanistic basis for understanding rhythmic perception through surround modulation circuits. This opens new possibilities for understanding rhythmic perception based on the physiology of visual processing. For example, human psychophysics work has shown that surround modulation is also tuned to spatial frequency and may not operate in the visual fovea (Petrov et al., 2005). So far, behavioural rhythms have mostly been investigated using peripheral stimuli (e.g. Landau et al., 2015) or using centrally-presented stimuli that are large enough to extend to the periphery where surround modulation effects are strong (Re et al., 2023). The existence of perceptual rhythms that are based in surround modulation circuitry prompts a deeper re-evaluation of oscillatory brain function and rhythmic perception. Surround modulation serves a computational purpose that is motivated ecologically by the need to efficiently process correlated spatial structure in the visual input (Schwartz & Simoncelli, 2001). While previous studies have proposed that neural oscillations may subserve efficiency-related goals within neural codes (Chalk et al., 2016), it remains unknown whether oscillatory neural activity may also serve a beneficial computational role in encoding the joint spatio-temporal features of the environment.

5.2 Methods

Participants and Ethics

Three healthy adults (all authors, one female, aged 25–30 years) gave their informed written consent to participate in the study. The study was approved by The University of Queensland Human Research Ethics Committee. Simulations

were run to confirm an appropriate number of trials per participant. All participants had normal or corrected-to-normal vision.

Stimuli, Task, and Procedure

Stimuli were modelled after those used in Petrov et al. (2005). The probe target was a standard Gabor (spatial frequency: 2 cycles/dva, random phase) in which 1.5 periods were visible, presented at 5 dva to the left/right of fixation for 16.7 ms. The surround stimulus was an oriented grating of the same spatial frequency as the target, presented as an annulus for 50 ms at 10% contrast. The outer edge of the annulus was 3 dva in diameter, and the inner edge allowed a blank separation of 1 period of the Gabor (~0.5 dva) between the surround and the visible part of the probe stimulus. Central eye fixation was maintained using a small black dot (width: 0.15 dva) inside a white larger white dot (width: 0.3 dva). Stimuli were displayed on a grey background using a luminance calibrated LCD monitor (1920 x 1080 pixels, 120 Hz refresh rate, Dell). Viewing distance was maintained at 68 cm using a chin rest.

Each trial started with a variable fixation period of 700–900 ms followed by a surround 'reset' stimulus. The probe stimulus appeared with a random timing relative to the surround reset (-100–1000 ms). Concurrent with the probe stimulus onset, the white portion of the central fixation dot was changed to a blue colour to cue the relevant time for judging the stimulus. Participants performed a detection task ('yes'/'no') and the target was absent on 50% of all trials.

Within a session, the orientation of the probe and surround aligned on half of the trials or rotated 90°orthogonally on the remaining half. Orientations were also counterbalanced using either only the two cardinal or the two oblique orientations within a session (the choice of which set was also counterbalanced across sessions). Stimuli were presented on either the left or right side of fixation for the entire session. This location was counterbalanced and randomised between the sessions of each participant.

Participants performed 800 trials per session, for a total 8 sessions per participant. In addition to this were several practice sessions of the task prior to participation in the main experiment, which were used to ensure each observer was able to perform the task well and produce stable estimates of threshold contrast. Breaks were enforced after every four minutes of continuous trial performance.

An adaptive psychophysical procedure was used to determine a threshold level of contrast visibility at which the probe stimulus was presented. This was based

on a four-parameter psychometric function (threshold, slope, lapse, criterion) that estimated signal detection theory d', and threshold was defined as the contrast value that produced performance at $d' = 1.5$. Eighty trials were used at the start of each session to establish psychometric parameters using an adaptive particle filter method (DiMattina, 2015). These trials were removed from subsequent analyses.

Analyses

Time Series Estimation

Time series data (y) were estimated using a sliding Gaussian window that was centred on times between -100 ms and 1000 ms, relative to surround stimulus onset, with a sampling resolution of 512 Hz. This was implemented using the dot product between the response data in the $[N, 1]$ vector x and weight matrix W,

$$y = W^T x, \tag{5.1}$$

with W defined as $e^{\frac{-(x-t)}{2*\sigma^2}}$, using the $[1, N]$ vector, t, of probe-surround interstimulus intervals. The parameter σ controls the standard deviation of the Gaussian temporal window. While the sliding window filter is generally low-pass, the value of σ can influence the inclusion of high frequencies. To mitigate bias from parameter choices, we chose a robust strategy and used five values of σ logarithmically spaced between 5–50 ms, averaging the resulting time series before any following analyses. To estimate accuracy time series that can be interpreted as a proportion, accuracy time series were computed by using data coded first as 1 = accurate, -1 = inaccurate, performing the sliding window computation, and then rescaling the result back into coding such that 1 = accurate, 0 = inaccurate.

Signal Detection Analyses

Time series of signal detection theory metrics were estimated using the same Gaussian sliding window approach. First, the hit rate, H, and false alarm rate, F, time series were computed using the matrix, W, defined above, and the vector, x:

$$H = \frac{Wx^{11} + m}{Wx^{11} + Wx^{10} + 2m}, \qquad F = \frac{Wx^{01} + m}{Wx^{01} + Wx^{00} + 2m}. \tag{5.2}$$

Here, x is a size $[N, 1]$ indicator variable that returns the value 1 on each trial based on participants' responses given stimulus presence, such that x^{11} indexes hits

and x^{01} indexes false alarms (x^{01} are misses and x^{00} are correct rejections). This creates time series that reflect the local hit rate and false alarm rate in each Gaussian window. A small offset ($m = 0.005$) was added to mitigate potential biases from small bin numbers. Following this, signal detection metrics for sensitivity (d') and criterion (c) are estimated as,

$$d' = \phi^{-1}(H) - \phi^{-1}(F) \tag{5.3}$$

$$c = -\frac{1}{2}\left(\phi^{-1}(H) + \phi^{-1}(F)\right), \tag{5.4}$$

using ϕ^{-1} as the inverse of the normal integral function.

Frequency Domain Analyses

Time series data prior to 100 ms post-reset stimulus were removed from analyses to avoid contamination from analysing transient effects close to stimulus presentation, which may not be oscillatory (Landau & Fries, 2012). Data were then demeaned, multiplied with a Hann window, and converted into complex frequency spectra using a Fast Fourier Transform (zero-padded to 4,096 samples). Frequency-domain coefficients were averaged across participants in the complex domain (which effectively averages in phase and in amplitude) before converting to spectral power via squaring the magnitude of the resulting spectrum.

Statistics

A bootstrapping approach was used to create 500 resampled datasets, with stratified resampling to ensure data was equated within participants and matched for the number of trials that were present/absent and iso-/cross-oriented. The *specparam* algorithm (Donoghue et al., 2020) was used to fit the resulting group averaged power spectra and separate its aperiodic and oscillatory components. This was done on frequency data ranging between 1 and 45 Hz, using default parameters of the algorithm. Across resampled datasets, we quantified the probability of detecting an oscillation between 2 and 15 Hz as well as the distribution of detected peak frequencies. A wide low-frequency band was chosen to quantify potential frequency variability in the (low) theta and alpha bands.

Lateral Inhibition Network Model

Simulations were run using the standard equations defined in Wilson and Cowan (1972) to create time series of activations in a coupled population of excitatory

and inhibitory neurons. Parameters were adjusted arbitrarily to create a stable oscillation in the coupled system. In principle, a range of parameter settings can produce the effects shown in Figure 5.1, but we omit this from examination as the simulations reflect a proof-of-concept for the computational ideas about oscillating and orientation-tuned surround modulation. To simulate the neural populations receiving an inhibitory signal from the spatial surround, a step-function impulse (lasting 20 ms) was delivered to the inhibitory population activation and at randomly jittered timepoint. In theory, these inhibitory surround signals are orientation-tuned, which is operationalised by the strength of this impulse. We compared two cases on the orientation-tuning function that reflect full vs no surround inhibition (Figure 5.1b). Stimulus input was modelled as a weaker 5 ms step-impulse provided to the excitatory population. The stimulus-evoked activity was decoded from the model as the maximum activation level in the excitatory population activation within the simulation, in line with a detection task where neural signals would be thresholded to inform a detection response.

Chapter 6

General Discussion

The goal of this thesis was to investigate how neural oscillations shape perceptual processing in vision. In this chapter, I provide an overview of the experimental and modelling work that I have presented. The implications of these findings for future experiments and some open questions for the field of research are discussed.

6.1 Overview of Findings

My aim was to investigate how neural oscillations shape the computational processes that give rise to visual perception. A critical part of tackling this question is understanding how alpha oscillations shape neural responses to external stimulation. Therefore, the first empirical chapter (Chapter 2) presented an investigation of alpha's effect on the human contrast response function. Analysing the manner in which alpha oscillations modulate the contrast response function can shed light on the potential computational processes that are affected by pulsed-inhibition in the visual system. To this end, observers' performance was measured in both detection and spatial discrimination tasks that involved brief stimuli presented in the visual periphery, at varying contrast intensities. Performance in both tasks was then examined with a signal detection theory (SDT) model incorporating a nonlinear transducer function, which specifies how contrast is nonlinearly converted into the corresponding strength of the internal representation.

The results of Chapter 2 showed that internal representations are generally inhibited by prestimulus alpha oscillations. This corresponded to fewer 'stimulus present' judgements in the detection task and less veridical reports of the location of the stimulus presented in the discrimination task. The SDT modelling revealed

that these effects can be explained by a multiplicative response gain modulation of observers' perceptual sensitivity, across the full contrast range. The key support for this comes from the inferred modulation effects on the latent sensory responses. In the detection task, latent responses were also shown to be additively modulated, reflecting either a change in baseline sensory response activation or a decision threshold bias (discussed further in this chapter).

In addition to revealing modulation of the contrast response function, an important contribution of Chapter 2 is that it profiled how sensory inhibition covaries with alpha power and phase. This was made possible by using a novel linking function that combined prestimulus alpha power and phase into an estimate of pulsed-inhibition at the time of stimulus presentation. This model of pulsed-inhibition is often assumed to be the basis of alpha's perceptual effects (Mathewson et al., 2011; Mazaheri & Jensen, 2010), and even finds support in the joint modulation of perception by alpha power and phase (Fakche et al., 2022), but had not previously been modelled in an explicit functional form. The model inferred the pattern of inhibition required to account for the reduced perceptual performance (through inhibition of response gain of the contrast response function), indicating that sensory responses are subject to a continuous, tonic form of inhibition that increases with alpha power, on top of the cycle-by-cycle pulsed inhibition linked to each alpha oscillation. The results of Chapter 2 therefore offer important insights into the sensory computations that are affected by alpha oscillations (i.e. neural response gain & baseline modulation), as well as how the putative pulsed-inhibition basis of alpha shapes sensory responses in an unexpected manner.

Another important aspect of visual system function, beyond how it encodes the visual environment in the magnitude of its sensory responses, is that it contains sufficient activity to discriminate environmental features. Orientation is a visual feature that early visual neurons are particularly responsive to (e.g. Hubel and Wiesel, 1962). In Chapter 3, I asked how alpha oscillations shaped neural representations of orientation that are encoded in the population activity of early visual neurons. Here, participants were tasked with reproducing the orientation of brief threshold-level stimuli that were presented during periods of high or low alpha power. High alpha power produced greater variability in continuous estimates relative to ground truth stimuli, as well as a greater bias in estimates away from the cardinal vertical and horizontal orientations. The latter finding can be explained by the fact that weak orientation stimuli are, in general, repulsively biased relative to

the cardinal orientations that dominate natural visual environments (de Gardelle et al., 2010; Wei & Stocker, 2015).

Perceptual biases reflect an efficient coding strategy in which neural tuning is preferentially allocated to more common signals (Barlow, 1961; Ganguli & Simoncelli, 2014; Girshick et al., 2011; Wei & Stocker, 2017). Using a Bayesian observer model that incorporates efficient allocation of neural resources, I showed that the key differences between orientation estimates as a function of alpha power were that high alpha produced greater noise in sensory estimates and more frequent guessing. In addition, Chapter 3 revealed that subjective visibility of stimuli was modulated by the alpha phase and that stimuli were experienced as less visible during high alpha power. Taken together, these findings suggest how alpha-band inhibition alters the representational strength of continuous visual features encoded in early visual cortex, with the result that the perception of orientation becomes more variable due to representational uncertainty and weakened subjective perception.

In Chapter 4, I created an observer model that aimed to explain the empirical results of the previous chapters in terms of realistic neural computations that underlie visual information processing. The model implemented a population code of V1-like orientation-tuned neurons that output Poisson spiking activity. Using Bayesian decision theory, it was possible to simulate spiking activity in response to visual stimuli, like those used in Chapters 2 and 3, and then decode visual representations contained in the spike-based population code. An important assumption of the model is that it receives spike-rate inhibition, mimicking the alpha-band inhibition of spiking activity that is found in neurophysiological recordings (Bastos et al., 2015; Bollimunta et al., 2008; Dougherty et al., 2017; Haegens et al., 2011; Spaak et al., 2012; van Kerkoerle et al., 2014). The model implements the functions of biological neurons that are modelled at a more abstract level in Chapter 2, using the SDT-transducer model, and in Chapter 3, using the efficient Bayesian observer. Thus, the model provides a way to investigate the precise neural mechanisms that link putative pulsed-inhibition to perceptual changes across detection and orientation estimation tasks.

Using simulations of the neural observer model, I was able to evaluate the potential ways in which pulsed-inhibition might affect selective functional properties of the neural population code. This was first applied to the contrast response function in a simulated detection task. This revealed that inhibiting the baseline and response gain of sensory activity, together, could account for alpha's

effects on the psychometric function, matching the SDT modelling conclusions from Chapter 2. Moreover, the model indicates that the same effects can result from another neural mechanism, in which stimulus-evoked activity is divisively normalised by pooling the activity of neighbouring neurons in the population into a suppressive field. To date, alpha's effects on the contrast response function have been previously studied in the literature and attributed to response gain modulation (Chaumon & Busch, 2014). However, the suppressive field account of modulation is a novel hypothesis generated from the model. Fitting the model to human detection performance (from Chapter 2) supported either of these accounts acting in combination with baseline modulation, but the results were equivocal about which specific account had more empirical support.

The neural observer model was also analysed in the context of an orientation estimation task by efficiently reallocating the neural population's orientation tuning. This reallocation was undertaken using the Fisher Information from observers' behavioural reports in Chapter 3, which permits a unified description of representational capacities in both model and human systems. Inhibition of the spiking rate of the efficient neural population produced stronger repulsive cardinal biases in the orientation estimates that could be decoded from the spikes. Thus, the model offered computational verification that pulsed-inhibition of visual activity is sufficient to account for the perceptual effects of alpha that were demonstrated for biased orientation judgements in Chapter 3.

The final study (Chapter 5) investigated whether rhythmic visual processing arises from inhibitory surround modulation. In vision, surround modulation reflects a suppressive influence of spatial context that is mediated, in theory, by neural computations that resemble the suppressive fields that enter into divisive normalisation (featured in Chapter 4; Angelucci and Bressloff, 2006; Angelucci et al., 2017). I developed a behavioural experiment to examine the hypothesis that surround modulation gives rise to oscillatory perception, in line with a proposed pulsed-inhibition mechanism facilitated by inhibitory lateral connections between neural populations. Investigating this using behaviour alone was made possible by using a stimulus to align and reset the phase of unobserved neural oscillations that may be ongoing in the brain (cf. Kienitz et al., 2021). Observers were tasked with detecting a stimulus that was presented at various latencies relative to a high contrast oriented spatial surround (e.g. an oriented reset stimulus). Using SDT, it was found that perceptual sensitivity fluctuated at a ~5 Hz rhythm. Crucially, rhythmic perception depended on orientation alignment between the

centre and the surround stimulus, consistent with orientation-tuned suppressive effects in previous psychophysical studies (Petrov et al., 2005). This suggests that the circuits that implement surround modulation may be intrinsically rhythmic. Ultimately, this is consistent with the role of inhibitory interneurons that mediate lateral inhibition (Del Rosario et al., 2025; Ferguson & Cardin, 2020) and neural oscillations (Jones et al., 2000).

6.2 Implications for Future Research

Toward a Theory of Rhythmic Visual Perception

The results presented in this thesis have implications for a theory-based understanding of rhythmic perception. In particular, the neural observer model from Chapter 4 suggests a way to interrogate the range of perceptual effects that should be oscillatory given the assumptions of the model and thus to determine the sufficiency of those assumptions. The model I have developed assumes that the primary mechanistic action of alpha oscillations is spike rate inhibition. I have shown that this assumption is sufficient to explain the effects of pulsed-inhibition in both detection and orientation estimation tasks using a single model of the visual encoding processes. Importantly, this implies that the relationship between perceptual behaviour and neural oscillations is determined by what specific visual information is encoded in the sensory activity that receives inhibition and how this is relevant to the behavioural task at hand.

At its core the neural observer model is a biologically inspired description of visual processing that can easily integrate the assumptions of pulsed-inhibition. This is a different approach from previous models that have been used to explore oscillations in perception. For example, the baseline sensory excitability model (BSEM) proposes a way to understand alpha's effects on visual perception using a SDT-based model of judgements about detection, discrimination, and confidence. Although this approach can describe sensory computations at an abstract explanatory level, it cannot offer a biologically plausible account of perceptual phenomena that are ultimately the result of fine-grained discriminability in neural representations (for example, the perceptual biases arising from efficient population coding, discussed in Chapters 3 and 4). This suggests that a simplistic notion of neural excitability (or, conversely, inhibition) may not capture the more realistic implications of fluctuations in the activity of complex visual populations.

Across the studies I have presented, a core finding is that neural oscillations interact with the canonical functions of visual neurons. Repulsive cardinal orientation biases (Chapter 3) show that alpha-based inhibition reduces the fidelity of perceptual representations. Importantly, these representations must be distributed among a neural population code that is efficiently adapted to environmental orientation statistics. Divisive normalisation is a feature of neural populations that implements both gain control and surround modulation. These features of population activity lie at the heart of the perceptual effects I investigated in the other chapters. For instance, divisive gain control affects how stimulus intensities are encoded in the activation strengths of individual neurons, which is central to the contrast-dependent effects of alpha on the psychometric functions in Chapter 2. Chapter 5 shows that surround modulation effects produce a theta rhythm in perceptual sensitivity, suggesting that inhibitory connections between neural populations mediate oscillating feature-tuned suppression. In general, these findings indicate that the perceptual impact of oscillations arises from inhibitory brain circuits that shape the content and nature of sensory information, rather than simply scaling the gain of the brain's sensory responses to the environment. The broader implication is that rhythmic visual perception should be viewed as an expression of how neural oscillations dynamically regulate the organisation and readout of sensory information.

The Effects of Alpha Oscillations on Signal Processing

An interesting point of divergence between the BSEM and the findings of the present thesis is the mechanism attributable to oscillatory modulation of sensory responses. A core question is whether these mechanisms alter the visual neurons' processing of input signals or reflect intrinsic noise variability that establishes a baseline level of activation. The BSEM supports the latter and posited that alpha oscillations impart an additive boost in sensory excitability that does not depend on the stimulus signal's intensity (Samaha et al., 2020). In contrast, the data presented in Chapter 2 support both an additive and a stimulus-dependent (multiplicative) change in the strength of the sensory response (which grows with the intensity of the stimulus). This raises the question: why did previous studies that informed the BSEM fail to show alpha-related modulation of sensitivity to sensory stimuli?

A reasonable explanation is that when both additive and multiplicative changes are present, this fact cannot be revealed by detection experiments that use

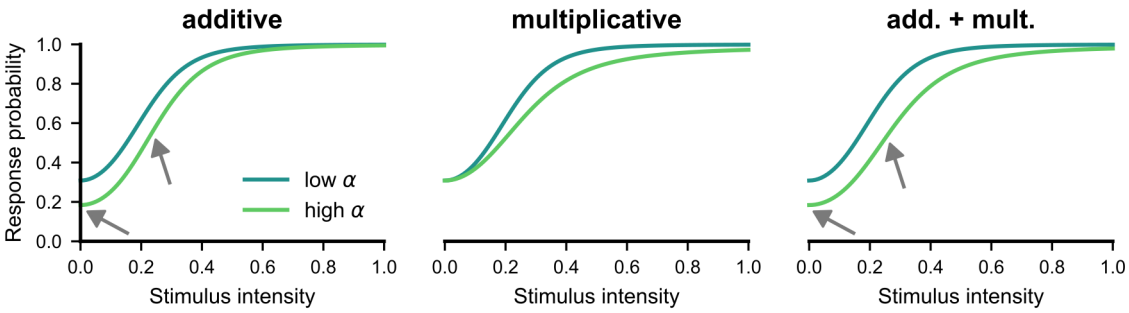

Figure 6.1: Additive and multiplicative changes in the psychometric function

The psychometric function in detection is simulated under conditions of high and low alpha power. Left shows the effect of modulating baseline sensory activity which corresponds to spontaneous neural activity that is not related to the stimulus (i.e. 'noise'). The grey arrows show the two points on the psychometric function that would be examined in a typical detection task where the stimulus is either absent or present (at threshold-level intensity). Middle shows the effect of modulating the response strength of stimulus driven activity. Right shows the combined effect.

only a single stimulus intensity level. For instance, a small shift in baseline sensory activity will increase the probability of detection when stimuli are not presented, increasing the false alarm rate. While this additive change will have reduced slightly at threshold, this is compensated for by the concomitant multiplicative effects that also operate in this contrast range. This is illustrated in Figure 6.1 which shows additive changes in sensory evidence, in isolation, shift the psychometric function vertically, and this is strongest at zero intensity and also apparent at threshold. A SDT analysis sees only two points on the psychometric function, and when the baseline and response gain effects coexist, the modulation at those points is consistent with a criterion change. However, the evidence for stimulus-dependent multiplicative changes in sensory evidence (i.e. d') is revealed only by nuances across the psychometric function, such as modulation at intensities above threshold. Ultimately, to understand how alpha oscillations shape the strength of sensory response, it is important to account for the nonlinearities involved in neural signal processing, such as the contrast transducer function that underpins the psychometric function.

The concept of perceptual *noise* is particularly relevant to understanding alpha's impacts on perceptual processing. Noise can have different meanings, such as (1) spiking activity that is unrelated to a meaningful signal, or (2) the uncertainty in perceptual representations or population codes. In the BSEM, the additive baseline shift reflects increased activation in sensory neurons that is not related to the incoming visual signals. However, this spontaneous activity can be facilitative in the context of a detection task. Consider, for example, the phenomenon of

stochastic resonance, where adding small amounts of noise can boost weak signals over a detection threshold (e.g. Linkenkaer-Hansen et al., 2004). In Chapter 3, I showed that the variability and biases in observers' perceptual reports can be explained by an increase in the noise in orientation representations, during high alpha power. In this case, noise in perceptual representations does not correspond to spontaneous activations of sensory neurons. The neural observer model in Chapter 4 revealed that the neural basis of uncertainty in perceptual estimation is, instead, the suppression of spiking responses. This means that, in contrast to detection, increased spontaneous activity does not facilitate the perceptual representation of orientation stimuli. Ultimately, this suggests that it is the signal-to-noise levels of task-relevant sensory activity that shape alpha's impacts on perception and behaviour. A more explicit way to study how alpha oscillations shape signal-to-noise levels in perception would be to use an external noise psychophysical paradigm (see Lu and Dosher, 2008 for review). This paradigm infers the levels of perceptual noise that are internal to the neural system by titrating the noise levels in an external stimulus. Such paradigms have not been applied in the context of oscillations, but could be used to characterise how endogenous neural states shape signal-to-noise processing in the visual system.

Extending the Neural Observer Model

A notable assumption of the neural observer model (Chapter 4) is that spiking activity is decoded in a fixed window of time. Extensions to this model could implement more realistic assumptions about how spiking activity is integrated into a decision-relevant signal over time. To start with, extending the model to explicitly account for the temporal evolution of spiking activity alongside pulsed-inhibition could yield a tool for understanding the joint influences of oscillatory phase and frequency on perception. One idea is that alpha oscillations underlie discrete clustering of perceptual information in time (VanRullen, 2016) consistent with evidence that alpha *frequency* modulates whether stimuli presented nearby in time are integrated into a single percept (more likely with low alpha frequency) or are perceived separately due to stimuli arriving at separate phases (more likely at higher frequency; Samaha and Romei, 2024).

An interesting idea that could be explored with the model is that the spiking integration window might vary with stimulus contrast. While neural responses are obviously weaker for lower contrast stimuli, it has been shown that they also appear with greater latency and are sustained over a longer period of time after

stimulus onset (Albrecht et al., 2002; Groen et al., 2022). This raises the possibility that the integration should not be treated as identical for all stimulus contrasts: lower contrast stimuli may accumulate more temporally extended spiking activations. There is a potentially important consequence of this for phasic perception: longer stimulus-evoked traces in the brain will overlap multiple cycles of pulsed-inhibition. This could be the reason why in Chapter 2 I found weaker phase effects, which are inconsistent with a naive pulsed-inhibition model. Potentially, lower-contrast stimuli are processed in a more temporally extended manner which makes the effects of phasic inhibition much less distinct. The covariation of contrast influences on stimulus-evoked dynamics with phasic perception could be an important finding that is accounted for in further modelling analyses.

A limitation of the neural approach to observer modelling is that it is computationally expensive to fit the models. This arises from the need to generate many simulations of spiking activity, making inferences about unknown parameters time-consuming and costly. I was able to overcome some of these challenges in Chapter 4 using simulation-focused analyses to examine the effects of different parameter settings *a priori* of the observed data. This limited the number of models that needed to be fit. Future applications may be able to solve this problem using simulation-based inference methods that train neural networks to perform likelihood-based inferences (Deistler et al., 2025).

Pulsed-Inhibition Beyond the Early Visual Cortex

The studies in this thesis have focused on explaining perception primarily through the lens of functions within early visual neurons. However, neural oscillations are found throughout the brain, and the impact of this is important to consider. Recall that, in Chapter 1, I reviewed evidence that alpha oscillations arise in the cortex through mechanisms that implement feedback from higher-level visual regions (van Kerkoerle et al., 2014). Dynamics like these are not explicit in the neural observer model I developed, but their influence could be examined in future work. A simple extension of the model could operationalise pulsed-inhibition as an inhibitory signal received from higher-order regions, consistent with neurophysiological models and theory (Bastos et al., 2012; Bastos et al., 2015). The surround modulation model presented in Chapter 5 already implements a form of lateral inhibition that controls the synchrony of spatially distributed neural populations. Such inhibitory surround interactions, if driven by hierarchically-higher brain regions, are a neurally plausible mechanism that could mediate phenomena

such as attentional sampling (Kienitz et al., 2018, 2021; Landau, 2018; Re et al., 2025).

Given that the dominant oscillatory activity in the visual cortex lies in the alpha-band (8–14 Hz), it is interesting that I found rhythmic surround modulation at a ~5 Hz theta-band rhythm in Chapter 5. Attentional rhythms are thought to operate a theta frequency (Landau et al., 2015) and result from alternating attentional sampling of alpha-frequency perceptual processing, effectively halving the frequency when attention is split between two objects or locations (Fiebelkorn & Kastner, 2019; Landau et al., 2015; Re et al., 2025). So, while perceptual processing might oscillate at an alpha rhythm, attention-related sampling between spatial locations could result in additional theta-frequency modulation of perception within a location. One possibility is that the theta-band surround modulation that I found reflects the influence of additional attention-related sampling, perhaps from the requirement to attend at fixation and also to the peripheral stimulus location. This suggests that the surround modulation effect could operate at an alpha frequency, but this is masked by the attentional demands of the task. Future work is needed to resolve this question of whether the rhythms induced by surround modulation are truly a theta-frequency effect or a product of attentional sampling.

Given the correlations between posterior alpha oscillations and attention (Peylo et al., 2021), it will be important for future work to understand how oscillations shape perceptual processing in the context of attention-related psychophysical effects. For instance, spatial attention modulates the contrast response function, as well as the feature-tuning of visual neurons (see Carrasco, 2011 for review). The normalisation model of attention (Reynolds & Heeger, 2009) is a successful account of the computational effects of visual attention, which is built on divisive normalisation computations. This was proposed to account for the varying effects of spatial attention on neural contrast response functions. It predicts that attended features receive either contrast or response gain boosts, which critically depend on the size of the processed stimulus relative to the spatial extent of attention. In the normalisation model, attention is operationalised as a boost in the gain of neural responses. This is compatible with response gain modulations from pulsed-inhibition, although it is framed in opposite terms as an increase. Previous modelling work has explored how high-frequency gamma oscillations may be a necessary component of hierarchical networks that implement attention with divisive normalisation, and this suggests that gamma oscillations play a specific role

in stabilising such networks by preserving the discriminability of signals (Montijn et al., 2012). Beyond this, it is unknown whether low-frequency oscillations, particularly alpha, are involved in the normalisation model of attention.

Why is the Brain Rhythmic?

An interesting feature of the neural observer model is that the effects on visual perception depend on spiking rate and not specifically on alpha oscillations. In principle, it is possible for oscillations of any frequency to modulate spike rate (Schalk, 2015). This raises the important question: why are alpha oscillations the dominant modulator of visual processing? One answer is that alpha arises from biophysical constraints of the circuits implementing the oscillations (e.g. synaptic time constants or intrinsic membrane properties). Alternatively, dynamical models of cortical circuits suggest that oscillations emerge from recurrent interactions among connected neurons (e.g. Bastiaens et al., 2025). These explanations are mechanistic in the sense that they reveal the processes that give rise to oscillations. However, the possibility that surround modulation circuits produce low-frequency oscillations (Chapter 5) suggests that there may be a more radical explanation that is motivated by a functional purpose.

The inhibitory connections that produce surround modulation serve a computational role to 'sparsen' neural coding and make encoded information more independent in the context of the highly correlated inputs of the visual system (Schwartz & Simoncelli, 2001). This functional principle suggests that alpha-band oscillations may not arise arbitrarily from biophysical constraints alone, but rather emerge because rhythmic inhibition naturally solves a problem posed by the structure of natural images. For example, natural visual scenes exhibit strong spatial correlations, such that nearby pixels have similar intensities and features tend to cluster (Field, 1987; Girshick et al., 2011). Early visual neurons, with their overlapping receptive fields, then face the challenge of redundant encoding in this correlated environment. To combat this, surround modulation reduces this redundancy and sharpens neural representations, by suppressing the response to contextually predictable stimulus characteristics Carandini and Heeger, 2012.

If surround modulation circuits are inherently rhythmic (as suggested by Chapter 5's experimental findings), then the alpha rhythm may be not merely an epiphenomenon of inhibitory neural networks, but rather a functional consequence of dynamic gain control mediated by lateral inhibition. In this view, alpha oscillations are rhythmic precisely because efficient coding in a spatially and

temporally correlated world requires periodic suppression of redundant information. Although a speculative possibility, this suggests that the alpha rhythm might carry some adaptive value and enable the visual system to balance its sensitivity to stimulus-driven information against the efficiency demands imposed by environmental correlations. In the spatial domain, the centre-surround spatial structure of natural images has been well characterised and motivates the general centre-surround organisation of visual neurons (Zhang et al., 2025). It might be possible to apply similar analyses to the joint spatio-temporal structure of natural image movies, to test whether these ideas about alpha generation have a basis in the statistics of environmental temporal frequency structure.

6.3 Concluding Remarks

A major goal in visual neuroscience is to understand how observable brain activity corresponds to ongoing processing of the sensory environment. Neural oscillations represent a major source of cortical variability and control the responsiveness of processes that unfold in the visual cortex. My thesis aimed to investigate the influence of neural oscillations on visual perception using a model-based approach to understanding visual function. Across the empirical studies presented, I quantified how neural oscillations shape perceptual behaviours like the detection, discrimination, and estimation of visual features. I found compelling support for perceptual processing being modulated by a common mechanism that reflects pulsed-inhibition to the sensory response in early visual population codes. Collectively, my results suggest that neural oscillations act on canonical visual circuits that implement divisive normalisation and surround modulation in ways that dynamically sculpt our perception of the world.

References

Abril-Pla, O., Andreani, V., Carroll, C., Dong, L., Fonnesbeck, C. J., Kochurov, M., Kumar, R., Lao, J., Luhmann, C. C., Martin, O. A., Osthege, M., Vieira, R., Wiecki, T., & Zinkov, R. (2023). PyMC: A modern, and comprehensive probabilistic programming framework in Python. *PeerJ Computer Science*, *9*, e1516. https://doi.org/10.7717/peerj-cs.1516

Adelson, E. H., & Bergen, J. R. (1985). Spatiotemporal energy models for the perception of motion. *Journal of the Optical Society of America A*, *2*(2), 284. https://doi.org/10.1364/JOSAA.2.000284

Adelson, E. H., & Bergen, J. R. (1991, October 24). The Plenoptic Function and the Elements of Early Vision. In M. Landy & J. A. Movshon (Eds.), *Computational Models of Visual Processing*. The MIT Press. https://doi.org/10.7551/mitpress/2002.003.0004

Adrian, E. D., & Matthews, B. H. C. (1934). The berger rhythm: Potential changes from the occipital lobes in man. *Brain*, *57*(4), 355–385. https://doi.org/10.1093/brain/57.4.355

Akam, T., & Kullmann, D. M. (2010). Oscillations and Filtering Networks Support Flexible Routing of Information. *Neuron*, *67*(2), 308–320. https://doi.org/10.1016/j.neuron.2010.06.019

Akam, T., & Kullmann, D. M. (2014). Oscillatory multiplexing of population codes for selective communication in the mammalian brain. *Nature Reviews Neuroscience*, *15*(2), 111–122. https://doi.org/10.1038/nrn3668

Albrecht, D. G., & Hamilton, D. B. (1982). Striate cortex of monkey and cat: Contrast response function. *Journal of Neurophysiology*, *48*(1), 217–237. https://doi.org/10.1152/jn.1982.48.1.217

Albrecht, D. G., Geisler, W. S., Frazor, R. A., & Crane, A. M. (2002). Visual Cortex Neurons of Monkeys and Cats: Temporal Dynamics of the Contrast Response Function. *Journal of Neurophysiology*, *88*(2), 888–913. https://doi.org/10.1152/jn.2002.88.2.888

Alexander, K. E., Estepp, J. R., & Elbasiouny, S. M. (2020). Effects of Neuronic Shutter Observed in the EEG Alpha Rhythm. *eNeuro*, *7*(5). https://doi.org/10.1523/ENEURO.0171-20.2020

Alonso, J.-M., Usrey, W. M., & Reid, R. C. (2001). Rules of Connectivity between Geniculate Cells and Simple Cells in Cat Primary Visual Cortex. *Journal of Neuroscience*, *21*(11), 4002–4015. https://doi.org/10.1523/JNEUROSCI.21-11-04002.2001

Angelucci, A., Bijanzadeh, M., Nurminen, L., Federer, F., Merlin, S., & Bressloff, P. C. (2017). Circuits and mechanisms for surround modulation in visual cortex. *Annual Review of Neuroscience*, *40*(1), 425–451. https://doi.org/10.1146/annurev-neuro-072116-031418

Angelucci, A., & Bressloff, P. C. (2006, January 1). Contribution of feedforward, lateral and feedback connections to the classical receptive field center and extra-classical receptive field surround of primate V1 neurons. In S. Martinez-Conde, S. L. Macknik, L. M. Martinez, J. .-. Alonso, & P. U. Tse (Eds.), *Progress in Brain Research* (pp. 93–120, Vol. 154). Elsevier. https://doi.org/10.1016/S0079-6123(06)54005-1

Anzai, A., Peng, X., & Van Essen, D. C. (2007). Neurons in monkey visual area V2 encode combinations of orientations. *Nature Neuroscience*, *10*(10), 1313–1321. https://doi.org/10.1038/nn1975

Balaban, H., & Ullman, T. D. (2025). Physics versus graphics as an organizing dichotomy in cognition. *Trends in Cognitive Sciences*, *29*(11), 985–996. https://doi.org/10.1016/j.tics.2025.05.003

Balestrieri, E., & Busch, N. A. (2022). Spontaneous alpha-band oscillations bias subjective contrast perception. *Journal of Neuroscience*, *42*(25), 5058–5069. https://doi.org/10.1523/JNEUROSCI.1972-21.2022

Barlow, H. B. (1961). Possible Principles Underlying the Transformations of Sensory Messages. In W. A. Rosenblith (Ed.), *Sensory Communication* (pp. 216–234). The MIT Press. https://doi.org/10.7551/mitpress/9780262518420.003.0013

Barlow, H. B. (1972). Single Units and Sensation: A Neuron Doctrine for Perceptual Psychology? *Perception*, *1*(4), 371–394. https://doi.org/10.1068/p010371

Bastiaens, S. P., Momi, D., & Griffiths, J. D. (2025). A comprehensive investigation of intracortical and corticothalamic models of the alpha rhythm. *PLOS Computational Biology*, *21*(4), e1012926. https://doi.org/10.1371/journal.pcbi.1012926

Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., & Friston, K. J. (2012). Canonical Microcircuits for Predictive Coding. *Neuron*, *76*(4), 695–711. https://doi.org/10.1016/j.neuron.2012.10.038 WOS:000311977900005.

Bastos, A. M., Vezoli, J., Bosman, C. A., Schoffelen, J.-M., Oostenveld, R., Dowdall, J. R., De Weerd, P., Kennedy, H., & Fries, P. (2015). Visual Areas Exert Feedforward and Feedback Influences through Distinct Frequency Channels. *Neuron*, *85*(2), 390–401. https://doi.org/10.1016/j.neuron.2014.12.018

Baumgarten, T. J., Neugebauer, J., Oeltzschner, G., Füllenbach, N.-D., Kircheis, G., Häussinger, D., Lange, J., Wittsack, H.-J., Butz, M., & Schnitzler, A. (2018). Connecting occipital alpha band peak frequency, visual temporal resolution, and occipital GABA levels in healthy participants and hepatic encephalopathy patients. *NeuroImage: Clinical*, *20*, 347–356. https://doi.org/10.1016/j.nicl.2018.08.013

Beale, H., Mattingley, J., & Harris, A. (2026, February 20). *A joint alpha power–phase dynamic shapes visual sensitivity*. https://doi.org/10.64898/2026.02.19.706926

Beck, J. M., Latham, P. E., & Pouget, A. (2011). Marginalization in Neural Circuits with Divisive Normalization. *The Journal of Neuroscience*, *31*(43), 15310–15319. https://doi.org/10.1523/JNEUROSCI.1706-11.2011

Benedetto, A., & Morrone, M. C. (2019). Visual sensitivity and bias oscillate phase-locked to saccadic eye movements. *Journal of Vision*, *19*(14), 15. https://doi.org/10.1167/19.14.15

Benedetto, A., Spinelli, D., & Morrone, M. C. (2016). Rhythmic modulation of visual contrast discrimination triggered by action. *Proceedings of the Royal Society B: Biological Sciences*, *283*(1831), 20160692. https://doi.org/10.1098/rspb.2016.0692

Benwell, C. S. Y., Coldea, A., Harvey, M., & Thut, G. (2022). Low pre-stimulus EEG alpha power amplifies visual awareness but not visual sensitivity. *European Journal of Neuroscience*, *25*(11–12). https://doi.org/10.1111/ejn.15166

Benwell, C. S. Y., Tagliabue, C. F., Veniero, D., Cecere, R., Savazzi, S., & Thut, G. (2017). Prestimulus EEG Power Predicts Conscious Awareness But Not Objective Visual Performance. *eneuro*, *4*(6), ENEURO.0182–17.2017. https://doi.org/10.1523/ENEURO.0182-17.2017

Berger, H. (1929). Über das Elektrenkephalogramm des Menschen. *Archiv für Psychiatrie und Nervenkrankheiten*, *87*(1), 527–570. https://doi.org/10.1007/BF01797193

Bollimunta, A., Chen, Y., Schroeder, C. E., & Ding, M. (2008). Neuronal Mechanisms of Cortical Alpha Oscillations in Awake-Behaving Macaques. *Journal of Neuroscience*, *28*(40), 9976–9988. https://doi.org/10.1523/JNEUROSCI.2699-08.2008

Bollimunta, A., Mo, J., Schroeder, C. E., & Ding, M. (2011). Neuronal Mechanisms and Attentional Modulation of Corticothalamic Alpha Oscillations. *Journal of Neuroscience*, *31*(13), 4935–4943. https://doi.org/10.1523/JNEUROSCI.5580-10.2011

Bonnefond, M., Jensen, O., & Clausner, T. (2024). Visual Processing by Hierarchical and Dynamic Multiplexing. *eNeuro*, *11*(11). https://doi.org/10.1523/ENEURO.0282-24.2024

Bonnefond, M., Kastner, S., & Jensen, O. (2017). Communication between Brain Areas Based on Nested Oscillations. *eNeuro*, *4*(2), ENEURO.0153–16.2017. https://doi.org/10.1523/ENEURO.0153-16.2017

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*(4), 433–436. https://doi.org/10.1163/156856897X00357

Briggs, F. (2017, February 27). Mammalian Visual System Organization. In S. M. Sherman (Ed.), *Oxford Research Encyclopedia of Neuroscience* (1st ed.). Oxford University PressNew York, NY. https://doi.org/10.1093/acrefore/9780190264086.013.66

Briggs, F. (2020). Role of Feedback Connections in Central Visual Processing. *Annual Review of Vision Science*, *6*(1), 313–334. https://doi.org/10.1146/annurev-vision-121219-081716

Brookshire, G. (2022). Putative rhythms in attentional switching can be explained by aperiodic temporal structure. *Nature Human Behaviour*, 1–12. https://doi.org/10.1038/s41562-022-01364-0

Brunel, N., & Nadal, J.-P. (1998). Mutual Information, Fisher Information, and Population Coding. *Neural Computation*, *10*(7), 1731–1757. https://doi.org/10.1162/089976698300017115

Buergers, S., & Noppeney, U. (2022). The role of alpha oscillations in temporal binding within and across the senses. *Nature Human Behaviour*. https://doi.org/10.1038/s41562-022-01294-x

Buffalo, E. A., Fries, P., Landman, R., Buschman, T. J., & Desimone, R. (2011). Laminar differences in gamma and alpha coherence in the ventral stream. *Proceedings of the National Academy of Sciences*, *108*(27), 11262–11267. https://doi.org/10.1073/pnas.1011284108

Busch, N. A., Dubois, J., & VanRullen, R. (2009). The phase of ongoing EEG oscillations predicts visual perception. *Journal of Neuroscience*, *29*(24), 7869–7876. https://doi.org/10.1523/JNEUROSCI.0113-09.2009

Busch, N. A., & VanRullen, R. (2010). Spontaneous EEG oscillations reveal periodic sampling of visual attention. *Proceedings of the National Academy of Sciences*, *107*(37), 16048–16053. https://doi.org/10.1073/pnas.1004801107

Busse, L., Wade, A. R., & Carandini, M. (2009). Representation of Concurrent Stimuli by Population Activity in Visual Cortex. *Neuron*, *64*(6), 931–942. https://doi.org/10.1016/j.neuron.2009.11.004

Buzsáki, G., Anastassiou, C. A., & Koch, C. (2012). The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience*, *13*(6), 407–420. https://doi.org/10.1038/nrn3241

Buzsáki, G., Logothetis, N., & Singer, W. (2013). Scaling Brain Size, Keeping Timing: Evolutionary Preservation of Brain Rhythms. *Neuron*, *80*(3), 751–764. https://doi.org/10.1016/j.neuron.2013.10.002

Callaway, E., & Alexander, J. D. (1960). The Temporal Coding of Sensory Data: An Investigation of Two Theories. *The Journal of General Psychology*, *62*(2), 293–309. https://doi.org/10.1080/00221309.1960.9920419

Callaway, E., & Yeager, C. L. (1960). Relationship between Reaction Time and Electroencephalographic Alpha Phase. *Science*, *132*(3441), 1765–1766. https://doi.org/10.1126/science.132.3441.1765

Canavier, C. C. (2015). Phase-resetting as a tool of information transmission. *Current Opinion in Neurobiology*, *31*, 206–213. https://doi.org/10.1016/j.conb.2014.12.003

Carandini, M. (2004). Receptive fields and suppressive fields in the early visual system. In M. Gazzaniga (Ed.), *Cognitive Neurosciences* (3rd ed., pp. 313–326). MIT Press.

Carandini, M., Demb, J. B., Mante, V., Tolhurst, D. J., Dan, Y., Olshausen, B. A., Gallant, J. L., & Rust, N. C. (2005). Do We Know What the Early Visual System Does? *The Journal of Neuroscience*, *25*(46), 10577–10597. https://doi.org/10.1523/JNEUROSCI.3726-05.2005

Carandini, M., & Heeger, D. J. (2012). Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, *13*(1), 51–62. https://doi.org/10.1038/nrn3136

Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research*, *51*(13), 1484–1525. https://doi.org/10.1016/j.visres.2011.04.012

Chalk, M., Gutkin, B., & Denève, S. (2016). Neural oscillations as a signature of efficient coding in the presence of synaptic delays (P. Latham, Ed.). *eLife*, *5*, e13824. https://doi.org/10.7554/eLife.13824

Chapeton, J., Haque, R., Wittig, J., Jr., Inati, S., & Zaghloul, K. (2019). Large-Scale Communication in the Human Brain Is Rhythmically Modulated through Alpha Coherence. *Current Biology*, *29*(17), 2801–2811.e5. https://doi.org/10.1016/j.cub.2019.07.014

Chaumon, M., & Busch, N. A. (2014). Prestimulus neural oscillations inhibit visual perception via modulation of response gain. *Journal of Cognitive Neuroscience*, *26*(11), 2514–2529. https://doi.org/10.1162/jocn_a_00653

Chen, L., Wu, B., Qiao, C., & Liu, D.-Q. (2020). Resting EEG in alpha band predicts individual differences in visual size perception. *Brain and Cognition*, *145*, 105625. https://doi.org/10.1016/j.bandc.2020.105625

Clayton, M. S., Yeung, N., & Kadosh, R. C. (2018). The many characters of visual alpha oscillations. *European Journal of Neuroscience*, *48*(7), 2498–2508. https://doi.org/10.1111/ejn.13747

Cole, S. R., & Voytek, B. (2017). Brain Oscillations and the Importance of Waveform Shape. *Trends in Cognitive Sciences*, *21*(2), 137–149. https://doi.org/10.1016/j.tics.2016.12.008

Conway, B. R., & Livingstone, M. S. (2006). Spatial and Temporal Properties of Cone Signals in Alert Macaque Primary Visual Cortex. *Journal of Neuroscience*, *26*(42), 10826–10846. https://doi.org/10.1523/JNEUROSCI.2091-06.2006

Cumming, B. G., & Parker, A. J. (1997). Responses of primary visual cortical neurons to binocular disparity without depth perception. *Nature*, *389*(6648), 280–283. https://doi.org/10.1038/38487

Davis, Z. W., Muller, L., & Reynolds, J. H. (2022). Spontaneous Spiking Is Governed by Broadband Fluctuations. *Journal of Neuroscience*, *42*(26), 5159–5172. https://doi.org/10.1523/JNEUROSCI.1899-21.2022

de Gardelle, V., Kouider, S., & Sackur, J. (2010). An oblique illusion modulated by visibility: Non-monotonic sensory integration in orientation processing. *Journal of Vision*, *10*(10), 6. https://doi.org/10.1167/10.10.6

DeCarlo, L. T. (1998). Signal detection theory and generalized linear models. *Psychological Methods*, *3*(2), 186–205. https://doi.org/10.1037/1082-989X.3.2.186

Deistler, M., Boelts, J., Steinbach, P., Moss, G., Moreau, T., Gloeckler, M., Rodrigues, P. L. C., Linhart, J., Lappalainen, J. K., Miller, B. K., Gonçalves, P. J., Lueckmann, J.-M., Schröder, C., & Macke, J. H. (2025, August 18). *Simulation-Based Inference: A Practical Guide*. arXiv: 2508.12939 [stat]. https://doi.org/10.48550/arXiv.2508.12939

Del Rosario, J., Coletta, S., Kim, S. H., Mobille, Z., Peelman, K., Williams, B., Otsuki, A. J., Del Castillo Valerio, A., Worden, K., Blanpain, L. T., Lovell, L., Choi, H., & Haider, B. (2025). Lateral inhibition in V1 controls neural and perceptual contrast sensitivity. *Nature Neuroscience*, 1–12. https://doi.org/10.1038/s41593-025-01888-4

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

DiMattina, C. (2015). Fast adaptive estimation of multidimensional psychometric functions. *Journal of Vision*, *15*(9), 5. https://doi.org/10.1167/15.9.5

Dong, D. W., & Atick, J. J. (1995). Statistics of natural time-varying images. *Network: Computation in Neural Systems*, *6*(3), 345. https://doi.org/10.1088/0954-898X/6/3/003

Donoghue, T., Haller, M., Peterson, E. J., Varma, P., Sebastian, P., Gao, R., Noto, T., Lara, A. H., Wallis, J. D., Knight, R. T., Shestyuk, A., & Voytek, B. (2020). Parameterizing neural power spectra into periodic and aperiodic components. *Nature Neuroscience*, *23*(12), 1655–1665. https://doi.org/10.1038/s41593-020-00744-x

Dou, W., Morrow, A., Iemi, L., & Samaha, J. (2022). Pre-Stimulus Alpha-Band Phase Gates Early Visual Cortex Responses. *NeuroImage*, 119060. https://doi.org/10.1016/j.neuroimage.2022.119060

Dougherty, K., Cox, M. A., Ninomiya, T., Leopold, D. A., & Maier, A. (2017). Ongoing Alpha Activity in V1 Regulates Visually Driven Spiking Responses. *Cerebral Cortex*, *27*(2), 1113–1124. https://doi.org/10.1093/cercor/bhv304

Dugué, L., Marque, P., & VanRullen, R. (2011). The phase of ongoing oscillations mediates the causal relation between brain excitation and visual perception. *Journal of Neuroscience*, *31*(33), 11889–11893. https://doi.org/10.1523/JNEUROSCI.1161-11.2011

Dustman, R. E., & Beck, E. C. (1965). Phase of alpha brain waves, reaction time and visually evoked potentials. *Electroencephalography and Clinical Neurophysiology*, *18*(5), 433–440. https://doi.org/10.1016/0013-4694(65)90123-9

Ellingson, R. J. (1956). Brain waves and problems of psychology. *Psychological Bulletin*, *53*(1), 1–34. https://doi.org/10.1037/h0042562

Enroth-Cugell, C., & Robson, J. G. (1966). The contrast sensitivity of retinal ganglion cells of the cat. *The Journal of Physiology*, *187*(3), 517–552. https://doi.org/10.1113/jphysiol.1966.sp008107

Ergenoglu, T., Demiralp, T., Bayraktaroglu, Z., Ergen, M., Beydagi, H., & Uresin, Y. (2004). Alpha rhythm of the EEG modulates visual detection performance in humans. *Cognitive Brain Research*, *20*(3), 376–383. https://doi.org/10.1016/j.cogbrainres.2004.03.009

Fakche, C., VanRullen, R., Marque, P., & Dugué, L. (2022). A Phase-Amplitude Tradeoffs Predict Visual Perception. *eNeuro*, *9*(1). https://doi.org/10.1523/ENEURO.0244-21.2022

Fechner, G. T. (with Francis A. Countway Library of Medicine). (1860). *Elemente der psychophysik*. Leipzig : Breitkopf und Härtel. Retrieved June 15, 2026, from http://archive.org/details/elementederpsych001fech

Ferguson, K. A., & Cardin, J. A. (2020). Mechanisms underlying gain modulation in the cortex. *Nature Reviews Neuroscience*, *21*(2), 80–92. https://doi.org/10.1038/s41583-019-0253-y

Fiebelkorn, I. C., & Kastner, S. (2019). A Rhythmic Theory of Attention. *Trends in Cognitive Sciences*, *23*(2), 87–101. https://doi.org/10.1016/j.tics.2018.11.009

Fiebelkorn, I. C., Saalmann, Y. B., & Kastner, S. (2013). Rhythmic Sampling within and between Objects despite Sustained Attention at a Cued Location. *Current Biology*, *23*(24), 2553–2558. https://doi.org/10.1016/j.cub.2013.10.063

Field, D. J. (1987). Relations between the statistics of natural images and the response properties of cortical cells. *JOSA A*, *4*(12), 2379–2394. https://doi.org/10.1364/JOSAA.4.002379

Fitzpatrick, D. (2000). Seeing beyond the receptive field in primary visual cortex. *Current Opinion in Neurobiology*, *10*(4), 438–443. https://doi.org/10.1016/S0959-4388(00)00113-6

Fries, P. (2005). A mechanism for cognitive dynamics: Neuronal communication through neuronal coherence. *Trends in Cognitive Sciences*, *9*(10), 474–480. https://doi.org/10.1016/j.tics.2005.08.011

Fries, P. (2015). Rhythms for Cognition: Communication through Coherence. *Neuron*, *88*(1), 220–235. https://doi.org/10.1016/j.neuron.2015.09.034

Gabry, J., Simpson, D., Vehtari, A., Betancourt, M., & Gelman, A. (2019). Visualization in Bayesian workflow. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, *182*(2), 389–402. https://doi.org/10.1111/rssa.12378

Ganguli, D., & Simoncelli, E. P. (2014). Efficient sensory encoding and Bayesian inference with heterogeneous neural populations. *Neural Computation*, *26*(10), 2103–2134. https://doi.org/10.1162/NECO_a_00638

Ganguli, D., & Simoncelli, E. P. (2016). *Neural and perceptual signatures of efficient sensory coding*. arXiv: 1603.00058 [q-bio]. https://doi.org/10.48550/arXiv.1603.00058

García-Pérez, M. A., & Alcalá-Quintana, R. (2007). The transducer model for contrast detection and discrimination: Formal relations, implications, and an empirical test. *Spatial vision*, *20*(1–2), 5–43. https://doi.org/10.1163/156856807779369724

Geisler, W. S. (2011). Contributions of ideal observer theory to vision research. *Vision Research*, *51*(7), 771–781. https://doi.org/10.1016/j.visres.2010.09.027

Gho, M., & Varela, F. J. (1988). A quantitative assessment of the dependency of the visual temporal frame upon the cortical rhythm. *Journal de physiologie*, *83*(2), 95–101.

Gips, B., van der Eerden, J. P. J. M., & Jensen, O. (2016). A biologically plausible mechanism for neuronal coding organized by the phase of alpha oscillations. *European Journal of Neuroscience*, *44*(4), 2147–2161. https://doi.org/10.1111/ejn.13318

Girshick, A. R., Landy, M. S., & Simoncelli, E. P. (2011). Cardinal rules: Visual orientation perception reflects knowledge of environmental statistics. *Nature Neuroscience*, *14*(7), 926–932. https://doi.org/10.1038/nn.2831

Gold, J. I., & Shadlen, M. N. (2007). The Neural Basis of Decision Making. *Annual Review of Neuroscience*, *30*(1), 535–574. https://doi.org/10.1146/annurev.neuro.29.051605.113038

Groen, I. I. A., Piantoni, G., Montenegro, S., Flinker, A., Devore, S., Devinsky, O., Doyle, W., Dugan, P., Friedman, D., Ramsey, N. F., Petridou, N., & Winawer, J. (2022). Temporal Dynamics of Neural Responses in Human Visual Cortex. *Journal of Neuroscience*, *42*(40), 7562–7580. https://doi.org/10.1523/JNEUROSCI.1812-21.2022

Groppe, D. M., Makeig, S., & Kutas, M. (2009). Identifying reliable independent components via split-half comparisons. *NeuroImage*, *45*(4), 1199–1211. https://doi.org/10.1016/j.neuroimage.2008.12.038

Gur, M., Kagan, I., & Snodderly, D. M. (2005). Orientation and Direction Selectivity of Neurons in V1 of Alert Monkeys: Functional Relationships and Laminar Distributions. *Cerebral Cortex*, *15*(8), 1207–1221. https://doi.org/10.1093/cercor/bhi003

Haegens, S., Barczak, A., Musacchia, G., Lipton, M. L., Mehta, A. D., Lakatos, P., & Schroeder, C. E. (2015). Laminar Profile and Physiology of the α Rhythm in Primary Visual, Auditory, and Somatosensory Regions of Neocortex. *Journal of Neuroscience*, *35*(42), 14341–14352. https://doi.org/10.1523/JNEUROSCI.0600-15.2015

Haegens, S., Cousijn, H., Wallis, G., Harrison, P. J., & Nobre, A. C. (2014). Inter- and intra-individual variability in alpha peak frequency. *NeuroImage*, *92*, 46–55. https://doi.org/10.1016/j.neuroimage.2014.01.049

Haegens, S., Nácher, V., Luna, R., Romo, R., & Jensen, O. (2011). A-Oscillations in the monkey sensorimotor network influence discrimination performance by rhythmical inhibition of neuronal spiking. *Proceedings of the National Academy of Sciences*, *108*(48), 19377–19382. https://doi.org/10.1073/pnas.1117190108

Hahn, M., & Wei, X.-X. (2024). A unifying theory explains seemingly contradictory biases in perceptual estimation. *Nature Neuroscience*, 1–12. https://doi.org/10.1038/s41593-024-01574-x

Halgren, M., Ulbert, I., Bastuji, H., Fabó, D., Erőss, L., Rey, M., Devinsky, O., Doyle, W. K., Mak-McCully, R., Halgren, E., Wittner, L., Chauvel, P., Heit, G., Eskandar, E., Mandell, A., & Cash, S. S. (2019). The generation and propagation of the human alpha rhythm. *Proceedings of the National Academy of Sciences*, *116*(47), 23772–23782. https://doi.org/10.1073/pnas.1913092116

Hanslmayr, S., Aslan, A., Staudigl, T., Klimesch, W., Herrmann, C. S., & Bäuml, K.-H. (2007). Prestimulus oscillations predict visual perception performance between and within subjects. *NeuroImage*, *37*(4), 1465–1473. https://doi.org/10.1016/j.neuroimage.2007.07.011

Harris, A. M. (2023). Phase resets undermine measures of phase-dependent perception. *Trends in Cognitive Sciences*. https://doi.org/10.1016/j.tics.2022.12.008

Harris, A. M., & Beale, H. A. (2025, April 24). *Detecting behavioural oscillations with increased sensitivity: A modification of Brookshire's (2022) AR-surrogate method*. https://doi.org/10.7554/eLife.106141.1

Harris, A. M., Dux, P. E., & Mattingley, J. B. (2018). Detecting unattended stimuli depends on the phase of prestimulus neural oscillations. *The Journal of Neuroscience*, *38*(12), 3092–3101. https://doi.org/10.1523/JNEUROSCI.3006-17.2018

Harvey, B. M., Vansteensel, M. J., Ferrier, C. H., Petridou, N., Zuiderbaan, W., Aarnoutse, E. J., Bleichner, M. G., Dijkerman, H. C., van Zandvoort, M. J. E., Leijten, F. S. S., Ramsey, N. F., & Dumoulin, S. O. (2013). Frequency specific spatial interactions in human electrocorticography: V1 alpha oscillations reflect surround suppression. *NeuroImage*, *65*, 424–432. https://doi.org/10.1016/j.neuroimage.2012.10.020

Heeger, D. J. (1992). Normalization of cell responses in cat striate cortex. *Visual Neuroscience*, *9*(2), 181–197. https://doi.org/10.1017/S0952523800009640

Hermes, D., Petridou, N., Kay, K. N., & Winawer, J. (2019). An image-computable model for the stimulus selectivity of gamma oscillations (S. Haegens, Ed.). *eLife*, *8*, e47035. https://doi.org/10.7554/eLife.47035

Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology*, *160*(1), 106–154.2. Retrieved June 9, 2022, from https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1359523/

Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *The Journal of Physiology*, *195*(1), 215–243. https://doi.org/10.1113/jphysiol.1968.sp008455

Hughes, S., Lőrincz, M., Turmaine, M., & Crunelli, V. (2011). Thalamic Gap Junctions Control Local Neuronal Synchrony and Influence Macroscopic Oscillation Amplitude during EEG Alpha Rhythms. *Frontiers in Psychol-*

ogy, *2*. Retrieved March 15, 2023, from https://www.frontiersin.org/articles/10.3389/fpsyg.2011.00193

Hung, Y.-C., Schwartz, G., Cooper, E. A., & Alexander, E. (2025, July 10). *Homeostatic Adaptation of Optimal Population Codes under Metabolic Stress*. arXiv: 2507.07874 [cs]. https://doi.org/10.48550/arXiv.2507.07874

Iemi, L., & Busch, N. A. (2018). Moment-to-Moment Fluctuations in Neuronal Excitability Bias Subjective Perception Rather than Strategic Decision-Making. *eneuro*, *5*(3), ENEURO.0430–17.2018. https://doi.org/10.1523/ENEURO.0430-17.2018

Iemi, L., Busch, N. A., Laudini, A., Haegens, S., Samaha, J., Villringer, A., & Nikulin, V. V. (2019). Multiple mechanisms link prestimulus neural oscillations to sensory responses (S. Hanslmayr & L. Colgin, Eds.). *eLife*, *8*, e43620. https://doi.org/10.7554/eLife.43620

Iemi, L., Chaumon, M., Crouzet, S. M., & Busch, N. A. (2017). Spontaneous neural oscillations bias perception by modulating baseline excitability. *Journal of Neuroscience*, *37*(4), 807–819. https://doi.org/10.1523/JNEUROSCI.1432-16.2016

Iemi, L., Gwilliams, L., Samaha, J., Auksztulewicz, R., Cycowicz, Y. M., King, J.-R., Nikulin, V. V., Thesen, T., Doyle, W., Devinsky, O., Schroeder, C. E., Melloni, L., & Haegens, S. (2022). Ongoing neural oscillations influence behavior and sensory representations by suppressing neuronal excitability. *NeuroImage*, 118746. https://doi.org/10.1016/j.neuroimage.2021.118746

Jasper, H. H. (1936). Cortical Excitatory State and Variability in Human Brain Rhythms. *Science*, *83*(2150), 259–260. https://doi.org/10.1126/science.83.2150.259

Jastrow, J. (1892). Studies from the University of Wisconsin: On the Judgment of Angles and Positions of Lines. *The American Journal of Psychology*, *5*(2), 214–248. https://doi.org/10.2307/1410867

Jazayeri, M., & Movshon, J. A. (2006). Optimal representation of sensory information by neural populations. *Nature Neuroscience*, *9*(5), 690–696. https://doi.org/10.1038/nn1691
458 citations (Semantic Scholar/DOI) [2021-04-18].

Jensen, O. (2024). Distractor inhibition by alpha oscillations is controlled by an indirect mechanism governed by goal-relevant information. *Communications Psychology*, *2*(1), 1–11. https://doi.org/10.1038/s44271-024-00081-w

Jensen, O., & Bonnefond, M. (2026). The alpha rhythm: From physiology to behavior. *Physiological Reviews*, *106*(3), 1123–1159. https://doi.org/10.1152/physrev.00001.2025

Jensen, O., Bonnefond, M., & VanRullen, R. (2012). An oscillatory mechanism for prioritizing salient unattended stimuli. *Trends in Cognitive Sciences*, *16*(4), 200–206. https://doi.org/10.1016/j.tics.2012.03.002

Jensen, O., Gips, B., Bergmann, T. O., & Bonnefond, M. (2014). Temporal coding organized by coupled alpha and gamma oscillations prioritize visual processing. *Trends in Neurosciences*, *37*(7), 357–369. https://doi.org/10.1016/j.tins.2014.04.001
WOS:000339037900001.

Jensen, O., & Mazaheri, A. (2010). Shaping functional architecture by oscillatory alpha activity: Gating by inhibition. *Frontiers in Human Neuroscience*, *4*, 186. https://doi.org/10.3389/fnhum.2010.00186

Johnson, E. N., Hawken, M. J., & Shapley, R. (2001). The spatial transformation of color in the primary visual cortex of the macaque monkey. *Nature Neuroscience*, *4*(4), 409–416. https://doi.org/10.1038/86061

Jones, E. G. (2009). Synchrony in the Interconnected Circuitry of the Thalamus and Cerebral Cortex. *Annals of the New York Academy of Sciences*, *1157*(1), 10–23. https://doi.org/10.1111/j.1749-6632.2009.04534.x

Jones, S. R., Pinto, D. J., Kaper, T. J., & Kopell, N. (2000). Alpha-Frequency Rhythms Desynchronize over Long Cortical Distances: A Modeling Study. *Journal of Computational Neuroscience*, *9*(3), 271–291. https://doi.org/10.1023/A:1026539805445

Keitel, C., Ruzzoli, M., Dugué, L., Busch, N. A., & Benwell, C. S. Y. (2022). Rhythms in cognition: The evidence revisited. *European Journal of Neuroscience*, *55*(11–12), 2991–3009. https://doi.org/10.1111/ejn.15740

Kienitz, R., Schmid, M. C., & Dugué, L. (2021). Rhythmic sampling revisited: Experimental paradigms and neural mechanisms. *European Journal of Neuroscience*, *n/a*(n/a). https://doi.org/10.1111/ejn.15489

Kienitz, R., Schmiedt, J. T., Shapcott, K. A., Kouroupaki, K., Saunders, R. C., & Schmid, M. C. (2018). Theta Rhythmic Neuronal Activity and Reaction Times Arising from Cortical Receptive Field Interactions during Distributed Attention. *Current Biology*, *28*(15), 2377–2387.e5. https://doi.org/10.1016/j.cub.2018.05.086

Kleiner, M., Brainard, D., & Pelli, D. (2007). What's new in Psychtoolbox-3?, 89.

Klimesch, W. (2012). Alpha-band oscillations, attention, and controlled access to stored information. *Trends in Cognitive Sciences*, *16*(12), 606–617. https://doi.org/10.1016/j.tics.2012.10.007

Klimesch, W., Sauseng, P., & Hanslmayr, S. (2007). EEG alpha oscillations: The inhibition–timing hypothesis. *Brain Research Reviews*, *53*(1), 63–88. https://doi.org/10.1016/j.brainresrev.2006.06.003

Klímová, M., Bloem, I. M., & Ling, S. (2021). The specificity of orientation-tuned normalization within human early visual cortex. *Journal of Neurophysiology*, *126*(5), 1536–1546. https://doi.org/10.1152/jn.00203.2021

Knoblauch, K., & Maloney, L. T. (2012). *Modeling Psychophysical Data in R*. Springer New York. https://doi.org/10.1007/978-1-4614-4475-6

Kontsevich, L. L., & Tyler, C. W. (1999). Bayesian adaptive estimation of psychometric slope and threshold. *Vision Research*, *39*(16), 2729–2737. https://doi.org/10.1016/S0042-6989(98)00285-5

Kothe, C., Shirazi, S. Y., Stenner, T., Medine, D., Boulay, C., Grivich, M. I., Artoni, F., Mullen, T., Delorme, A., & Makeig, S. (2025). The lab streaming layer for synchronized multimodal recording. *Imaging Neuroscience*, *3*, IMAG.a.136. https://doi.org/10.1162/IMAG.a.136

Kusnir, F., & Landau, A. N. (2025). Temporality and the brain: The long and winding emergence of time in cognitive neuroscience. *Human Arenas*. https://doi.org/10.1007/s42087-025-00497-8

Landau, A. N. (2018). Neuroscience: A mechanism for rhythmic sampling in vision. *Current Biology*, *28*(15), R830–R832. https://doi.org/10.1016/j.cub.2018.05.081

Landau, A. N., & Fries, P. (2012). Attention Samples Stimuli Rhythmically. *Current Biology*, *22*(11), 1000–1004. https://doi.org/10.1016/j.cub.2012.03.054

Landau, A. N., Schreyer, H. M., van Pelt, S., & Fries, P. (2015). Distributed attention is implemented through theta-rhythmic gamma modulation. *Current Biology*, *25*(17), 2332–2337. https://doi.org/10.1016/j.cub.2015.07.048

Lansing, R. W. (1957). Relation of brain and tremor rhythms to visual reaction time. *Electroencephalography and Clinical Neurophysiology*, *9*(3), 497–504. https://doi.org/10.1016/0013-4694(57)90037-8

Lee, T.-W., Girolami, M., & Sejnowski, T. J. (1999). Independent Component Analysis Using an Extended Infomax Algorithm for Mixed Subgaussian

and Supergaussian Sources. *Neural Computation*, *11*(2), 417–441. https://doi.org/10.1162/089976699300016719

Lesmes, L., Lu, Z.-L., Baek, J., Tran, N., Dosher, B., & Albright, T. (2015). Developing Bayesian adaptive methods for estimating sensitivity thresholds (d′) in Yes-No and forced-choice tasks. *Frontiers in Psychology*, *6*. Retrieved July 21, 2022, from https://www.frontiersin.org/articles/10.3389/fpsyg.2015.01070

Li, B., Peterson, M. R., & Freeman, R. D. (2003). Oblique Effect: A Neural Basis in the Visual Cortex. *Journal of Neurophysiology*, *90*(1), 204–217. https://doi.org/10.1152/jn.00954.2002

Limbach, K., & Corballis, P. M. (2016). Prestimulus alpha power influences response criterion in a detection task. *Psychophysiology*, *53*(8), 1154–1164. https://doi.org/10.1111/psyp.12666

Lindsley, D. B. (1952). Psychological phenomena and the electroencephalogram. *Electroencephalography and Clinical Neurophysiology*, *4*(4), 443–456. https://doi.org/10.1016/0013-4694(52)90075-8

Linkenkaer-Hansen, K., Nikulin, V. V., Palva, S., Ilmoniemi, R. J., & Palva, J. M. (2004). Prestimulus Oscillations Enhance Psychophysical Performance in Humans. *Journal of Neuroscience*, *24*(45), 10186–10190. https://doi.org/10.1523/JNEUROSCI.2584-04.2004

Lopes da Silva, F. H., & Storm van Leeuwen, W. (1977). The cortical source of the alpha rhythm. *Neuroscience Letters*, *6*(2), 237–241. https://doi.org/10.1016/0304-3940(77)90024-6

Lopes da Silva, F. (2013). EEG and MEG: Relevance to Neuroscience. *Neuron*, *80*(5), 1112–1128. https://doi.org/10.1016/j.neuron.2013.10.017

Lopes da Silva, F. H. (1991). Neural mechanisms underlying brain waves: From neural membranes to networks. *Electroencephalography and Clinical Neurophysiology*, *79*(2), 81–93. https://doi.org/10.1016/0013-4694(91)90044-5

Lőrincz, M. L., Kékesi, K. A., Juhász, G., Crunelli, V., & Hughes, S. W. (2009). Temporal Framing of Thalamic Relay-Mode Firing by Phasic Inhibition during the Alpha Rhythm. *Neuron*, *63*(5), 683–696. https://doi.org/10.1016/j.neuron.2009.08.012

Lozano-Soldevilla, D., ter Huurne, N., Cools, R., & Jensen, O. (2014). GABAergic Modulation of Visual Gamma and Alpha Oscillations and Its Consequences for Working Memory Performance. *Current Biology*, *24*(24), 2878–2887. https://doi.org/10.1016/j.cub.2014.10.017

Lu, Z.-L., & Dosher, B. A. (2008). Characterizing observers using external noise and observer models: Assessing internal representations with external noise. *Psychological Review*, *115*(1), 44–82. https://doi.org/10.1037/0033-295X.115.1.44

Lundqvist, M., Bastos, A. M., & Miller, E. K. (2020). Preservation and Changes in Oscillatory Dynamics across the Cortical Hierarchy. *Journal of Cognitive Neuroscience*, *32*(10), 2024–2035. https://doi.org/10.1162/jocn_a_01600

Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience*, *9*(11), 1432–1438. https://doi.org/10.1038/nn1790

Ma, W. J., Kording, K. P., & Goldreich, D. (2023). *Bayesian models of perception and action: An introduction*. MIT Press.

Martinez-Trujillo, J. C., & Treue, S. (2004). Feature-based attention increases the selectivity of population responses in primate visual cortex. *Current Biology*, *14*(9), 744–751. https://doi.org/10.1016/j.cub.2004.04.028

Masquelier, T., Hugues, E., Deco, G., & Thorpe, S. J. (2009). Oscillations, Phase-of-Firing Coding, and Spike Timing-Dependent Plasticity: An Efficient Learning Scheme. *Journal of Neuroscience*, *29*(43), 13484–13493. https://doi.org/10.1523/JNEUROSCI.2207-09.2009

Mathewson, K. E., Gratton, G., Fabiani, M., Beck, D. M., & Ro, T. (2009). To See or Not to See: Prestimulus α Phase Predicts Visual Awareness. *Journal of Neuroscience*, *29*(9), 2725–2732. https://doi.org/10.1523/JNEUROSCI.3963-08.2009

Mathewson, K. E., Lleras, A., Beck, D. M., Fabiani, M., Ro, T., & Gratton, G. (2011). Pulsed out of awareness: EEG alpha oscillations represent a pulsed-inhibition of ongoing cortical processing. *Frontiers in Psychology*, *2*. https://doi.org/10.3389/fpsyg.2011.00099

Mazaheri, A., & Jensen, O. (2008). Asymmetric Amplitude Modulations of Brain Oscillations Generate Slow Evoked Responses. *The Journal of Neuroscience*, *28*(31), 7781–7787. https://doi.org/10.1523/JNEUROSCI.1631-08.2008

Mazaheri, A., & Jensen, O. (2010). Rhythmic Pulsing: Linking Ongoing Brain Activity with Evoked Responses. *Frontiers in Human Neuroscience*, *4*. https://doi.org/10.3389/fnhum.2010.00177

Meij, R. van der, Ede, F. van, & Maris, E. (2016). Rhythmic Components in Extracranial Brain Signals Reveal Multifaceted Task Modulation of Overlapping Neuronal Activity. *PLOS ONE*, *11*(6), e0154881. https://doi.org/10.1371/journal.pone.0154881

Melcón, M., Stern, E., Kessel, D., Arana, L., Poch, C., Campo, P., & Capilla, A. (2024). Perception of near-threshold visual stimuli is influenced by prestimulus alpha-band amplitude but not by alpha phase. *Psychophysiology*, *61*(5), e14525. https://doi.org/10.1111/psyp.14525

Mendoza-Halliday, D., Major, A. J., Lee, N., Lichtenfeld, M. J., Carlson, B., Mitchell, B., Meng, P. D., Xiong, Y. (, Westerberg, J. A., Jia, X., Johnston, K. D., Selvanayagam, J., Everling, S., Maier, A., Desimone, R., Miller, E. K., & Bastos, A. M. (2024). A ubiquitous spectrolaminar motif of local field potential power across the primate cortex. *Nature Neuroscience*, 1–14. https://doi.org/10.1038/s41593-023-01554-7

Michail, G., Toran Jenner, L., & Keil, J. (2021). Prestimulus alpha power but not phase influences visual discrimination of long-duration visual stimuli. *European Journal of Neuroscience*, *n/a*(n/a). https://doi.org/10.1111/ejn.15169

Michel, R., Dugué, L., & Busch, N. A. (2021). Distinct contributions of alpha and theta rhythms to perceptual and attentional sampling. *European Journal of Neuroscience*, *55*(11–12). https://doi.org/10.1111/ejn.15154

Milstein, V. (1974). Alpha wave phase and alpha attenuation. *Electroencephalography and Clinical Neurophysiology*, *37*(2), 167–172. https://doi.org/10.1016/0013-4694(74)90008-X

Montijn, J. S., Klink, P. C., & Van Wezel, R. J. A. (2012). Divisive Normalization and Neuronal Oscillations in a Single Hierarchical Framework of Selective Visual Attention. *Frontiers in Neural Circuits*, *6*. https://doi.org/10.3389/fncir.2012.00022

Morey, R. D. (2008). Confidence intervals from normalized data: A correction to cousineau (2005). *Tutorials in Quantitative Methods for Psychology*, *4*(2), 61–64. https://doi.org/10.20982/tqmp.04.2.p061

Morey, R. D., & Rouder, J. N. (2011). Bayes factor approaches for testing interval null hypotheses. *Psychological Methods*, *16*(4), 406–419. https://doi.org/10.1037/a0024377

Naka, K. I., & Rushton, W. a. H. (1966). S-potentials from colour units in the retina of fish (Cyprinidae). *The Journal of Physiology*, *185*(3), 536–555. https://doi.org/10.1113/jphysiol.1966.sp008001

Nikulin, V. V., Linkenkaer-Hansen, K., Nolte, G., Lemm, S., Müller, K. R., Ilmoniemi, R. J., & Curio, G. (2007). A novel mechanism for evoked responses in the human brain. *European Journal of Neuroscience*, *25*(10), 3146–3154. https://doi.org/10.1111/j.1460-9568.2007.05553.x

Nolan, H., Whelan, R., & Reilly, R. B. (2010). FASTER: Fully Automated Statistical Thresholding for EEG artifact Rejection. *Journal of Neuroscience Methods*, *192*(1), 152–162. https://doi.org/10.1016/j.jneumeth.2010.07.015

Nunez, P. L., & Srinivasan, R. (2006). *Electric fields of the brain: The neurophysics of EEG* (2nd ed). Oxford University Press.

Nunn, C. M. H., & Osselton, J. W. (1974). The Influence of the EEG Alpha Rhythm on the Perception of Visual Stimuli. *Psychophysiology*, *11*(3), 294–303. https://doi.org/10.1111/j.1469-8986.1974.tb00547.x

Olshausen, B. A., & Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, *381*(6583), 607–609. https://doi.org/10.1038/381607a0

Oostenveld, R., & Praamstra, P. (2001). The five percent electrode system for high-resolution EEG and ERP measurements. *Clinical Neurophysiology*, *112*(4), 713–719. https://doi.org/10.1016/S1388-2457(00)00527-7

Ostojic, S., & Brunel, N. (2011). From Spiking Neuron Models to Linear-Nonlinear Models. *PLOS Computational Biology*, *7*(1), e1001056. https://doi.org/10.1371/journal.pcbi.1001056

Padamsey, Z., Katsanevaki, D., Dupuy, N., & Rochefort, N. L. (2022). Neocortex saves energy by reducing coding precision during food scarcity. *Neuron*, *110*(2), 280–296.e10. https://doi.org/10.1016/j.neuron.2021.10.024

Pascucci, D., Menétrey, M. Q., Passarotto, E., Luo, J., Paramento, M., & Rubega, M. (2025). EEG brain waves and alpha rhythms: Past, current and future direction. *Neuroscience & Biobehavioral Reviews*, *176*, 106288. https://doi.org/10.1016/j.neubiorev.2025.106288

Pasupathy, A., & Connor, C. E. (2002). Population coding of shape in area V4. *Nature Neuroscience*, *5*(12), 1332–1338. https://doi.org/10.1038/972

Petrov, Y., Carandini, M., & McKee, S. (2005). Two Distinct Mechanisms of Suppression in Human Vision. *Journal of Neuroscience*, *25*(38), 8704–8707. https://doi.org/10.1523/JNEUROSCI.2871-05.2005

Peylo, C., Hilla, Y., & Sauseng, P. (2021). Cause or consequence? Alpha oscillations in visuospatial attention. *Trends in Neurosciences*, *44*(9), 705–713. https://doi.org/10.1016/j.tins.2021.05.004

Pilipenko, A., McGowan, A., & Samaha, J. (2026). Alpha-band phase modulates perceptual sensitivity by changing internal noise and sensory tuning (N. Faivre & H. Luo, Eds.). *eLife*, *15*, RP110000. https://doi.org/10.7554/eLife.110000

Pilipenko, A., & Samaha, J. (2024). Double Dissociation of Spontaneous Alpha-Band Activity and Pupil-Linked Arousal on Additive and Multiplicative Perceptual Gain. *Journal of Neuroscience*, *44*(19). https://doi.org/10.1523/JNEUROSCI.1944-23.2024

Pouget, A., Dayan, P., & Zemel, R. (2000). Information processing with population codes. *Nature Reviews Neuroscience*, *1*(2), 125–132. https://doi.org/10.1038/35039062

Pouget, A., Dayan, P., & Zemel, R. S. (2003). Inference and computation with population codes. *Annual Review of Neuroscience*, *26*, 381–410. https://doi.org/10.1146/annurev.neuro.26.041002.131112

Re, D., Inbar, M., Richter, C. G., & Landau, A. N. (2019). Feature-Based Attention Samples Stimuli Rhythmically. *Current Biology*, *29*(4), 693–699.e4. https://doi.org/10.1016/j.cub.2019.01.010

Re, D., Karvat, G., & Landau, A. N. (2023). Attentional Sampling between Eye Channels. *Journal of Cognitive Neuroscience*, *35*(8), 1350–1360. https://doi.org/10.1162/jocn_a_02018

Re, D., Kusnir, F., & Landau, A. N. (2025). Attentional sampling resolves competition along the visual hierarchy. *Trends in Cognitive Sciences*. https://doi.org/10.1016/j.tics.2025.06.004

Reynolds, J. H., & Heeger, D. J. (2009). The normalization model of attention. *Neuron*, *61*(2), 168–185. https://doi.org/10.1016/j.neuron.2009.01.002

Ringach, D. L. (2004). Mapping receptive fields in primary visual cortex. *The Journal of Physiology*, *558*(3), 717–728. https://doi.org/10.1113/jphysiol.2004.065771

Romei, V., Rihs, T., Brodbeck, V., & Thut, G. (2008). Resting electroencephalogram alpha-power over posterior sites indexes baseline visual cortex excitability. *NeuroReport*, *19*(2), 203. https://doi.org/10.1097/WNR.0b013e3282f454c4

Rust, N. C., Mante, V., Simoncelli, E. P., & Movshon, J. A. (2006). How MT cells analyze the motion of visual patterns. *Nature Neuroscience*, *9*(11), 1421–1431. https://doi.org/10.1038/nn1786

Ruzzoli, M., Torralba, M., Morís Fernández, L., & Soto-Faraco, S. (2019). The relevance of alpha phase in human perception. *Cortex*, *120*, 249–268. https://doi.org/10.1016/j.cortex.2019.05.012

Samaha, J., Iemi, L., Haegens, S., & Busch, N. A. (2020). Spontaneous brain oscillations and perceptual decision-making. *Trends in Cognitive Sciences*, *24*(8), 639–653. https://doi.org/10.1016/j.tics.2020.05.004

Samaha, J., Iemi, L., & Postle, B. R. (2017a). Prestimulus alpha-band power biases visual discrimination confidence, but not accuracy. *Consciousness and Cognition*, *54*, 47–55. https://doi.org/10.1016/j.concog.2017.02.005 WOS:000408790700005.

Samaha, J., Iemi, L., & Postle, B. R. (2017b). Prestimulus alpha-band power biases visual discrimination confidence, but not accuracy. *Consciousness and Cognition*, *54*, 47–55. https://doi.org/10.1016/j.concog.2017.02.005

Samaha, J., & Postle, B. R. (2015). The Speed of Alpha-Band Oscillations Predicts the Temporal Resolution of Visual Perception. *Current Biology*, *25*(22), 2985–2990. https://doi.org/10.1016/j.cub.2015.10.007

Samaha, J., & Romei, V. (2024). Alpha-Band Frequency and Temporal Windows in Perception: A Review and Living Meta-analysis of 27 Experiments (and Counting). *Journal of Cognitive Neuroscience*, *36*(4), 640–654. https://doi.org/10.1162/jocn_a_02069

Schalk, G. (2015). A general framework for dynamic cortical function: The function-through-biased-oscillations (FBO) hypothesis. *Frontiers in Human Neuroscience*, *9*. https://doi.org/10.3389/fnhum.2015.00352

Schalk, G., Marple, J., Knight, R. T., & Coon, W. G. (2017). Instantaneous voltage as an alternative to power- and phase-based interpretation of oscillatory brain activity. *NeuroImage*, *157*, 545–554. https://doi.org/10.1016/j.neuroimage.2017.06.014

Schneider, G. E. (1969). Two Visual Systems. *Science*, *163*(3870), 895–902. https://doi.org/10.1126/science.163.3870.895

Schoffelen, J.-M., Pesci, U. G., & Noppeney, U. (2024). Alpha Oscillations and Temporal Binding Windows in Perception—A Critical Review and Best Practice Guidelines. *Journal of Cognitive Neuroscience*, *36*(4), 655–690. https://doi.org/10.1162/jocn_a_02118

Schreglmann, S. R., Wang, D., Peach, R. L., Li, J., Zhang, X., Latorre, A., Rhodes, E., Panella, E., Cassara, A. M., Boyden, E. S., Barahona, M., Santaniello, S., Rothwell, J., Bhatia, K. P., & Grossman, N. (2021). Non-invasive

suppression of essential tremor via phase-locked disruption of its temporal coherence. *Nature Communications*, *12*(1), 363. https://doi.org/10.1038/s41467-020-20581-7

Schütt, H. H., & Wichmann, F. A. (2017). An image-computable psychophysical spatial vision model. *Journal of Vision*, *17*(12), 12. https://doi.org/10.1167/17.12.12

Schwartz, O., Hsu, A., & Dayan, P. (2007). Space and time in visual context. *Nature Reviews Neuroscience*, *8*(7), 522–535. https://doi.org/10.1038/nrn2155

Schwartz, O., & Simoncelli, E. P. (2001). Natural signal statistics and sensory gain control. *Nature Neuroscience*, *4*(8), 819–825. https://doi.org/10.1038/90526

Sclar, G., & Freeman, R. D. (1982). Orientation selectivity in the cat's striate cortex is invariant with stimulus contrast. *Experimental Brain Research*, *46*(3), 457–461. https://doi.org/10.1007/BF00238641

Sclar, G., Lennie, P., & DePriest, D. D. (1989). Contrast adaptation in striate cortex of macaque. *Vision Research*, *29*(7), 747–755. https://doi.org/10.1016/0042-6989(89)90087-4

Sclar, G., Maunsell, J. H. R., & Lennie, P. (1990). Coding of image contrast in central visual pathways of the macaque monkey. *Vision Research*, *30*(1), 1–10. https://doi.org/10.1016/0042-6989(90)90123-3

Silva, L. R., Amitai, Y., & Connors, B. W. (1991). Intrinsic Oscillations of Neocortex Generated by Layer 5 Pyramidal Neurons. *Science*, *251*(4992), 432–435. https://doi.org/10.1126/science.1824881

Silver, R. A. (2010). Neuronal arithmetic. *Nature Reviews Neuroscience*, *11*(7), 474–489. https://doi.org/10.1038/nrn2864

Smith, S. L. (1962). Angular estimation. *Journal of Applied Psychology*, *46*(4), 240–246. https://doi.org/10.1037/h0044445

Song, K., Meng, M., Chen, L., Zhou, K., & Luo, H. (2014). Behavioral Oscillations in Attention: Rhythmic α Pulses Mediated through θ Band. *Journal of Neuroscience*, *34*(14), 4837–4844. https://doi.org/10.1523/JNEUROSCI.4856-13.2014

Spaak, E., Bonnefond, M., Maier, A., Leopold, D. A., & Jensen, O. (2012). Layer-Specific Entrainment of Gamma-Band Neural Activity by the Alpha Rhythm in Monkey Visual Cortex. *Current Biology*, *22*(24), 2313–2318. https://doi.org/10.1016/j.cub.2012.10.020

Stevens, S. S. (1957). On the psychophysical law. *Psychological Review*, *64*(3), 153–181. https://doi.org/10.1037/h0046162

Tanner, W. P., & Swets, J. A. (1954). A decision-making theory of visual detection. *Psychological Review*, *61*(6), 401–409. https://doi.org/10.1037/h0058700

Taylor, R., & Bays, P. M. (2018). Efficient Coding in Visual Working Memory Accounts for Stimulus-Specific Variations in Recall. *Journal of Neuroscience*, *38*(32), 7132–7142. https://doi.org/10.1523/JNEUROSCI.1018-18.2018

Thut, G., Nietzel, A., Brandt, S. A., & Pascual-Leone, A. (2006). A-Band Electroencephalographic Activity over Occipital Cortex Indexes Visuospatial Attention Bias and Predicts Visual Target Detection. *Journal of Neuroscience*, *26*(37), 9494–9502. https://doi.org/10.1523/JNEUROSCI.0875-06.2006

Tomassini, A., Spinelli, D., Jacono, M., Sandini, G., & Morrone, M. C. (2015). Rhythmic Oscillations of Visual Contrast Sensitivity Synchronized with Action. *Journal of Neuroscience*, *35*(18), 7019–7029. https://doi.org/10.1523/JNEUROSCI.4568-14.2015

Tosato, T., Rohenkohl, G., Dowdall, J. R., & Fries, P. (2022). Quantifying rhythmicity in perceptual reports. *NeuroImage*, *262*, 119561. https://doi.org/10.1016/j.neuroimage.2022.119561

van Dijk, H., Schoffelen, J.-M., Oostenveld, R., & Jensen, O. (2008). Prestimulus Oscillatory Activity in the Alpha Band Predicts Visual Discrimination Ability. *Journal of Neuroscience*, *28*(8), 1816–1823. https://doi.org/10.1523/JNEUROSCI.1853-07.2008

van Hateren, J. H., & Ruderman, D. L. (1998). Independent component analysis of natural image sequences yields spatio-temporal filters similar to simple cells in primary visual cortex. *Proceedings of the Royal Society B: Biological Sciences*, *265*(1412), 2315–2320. https://doi.org/10.1098/rspb.1998.0577

van Kerkoerle, T., Self, M. W., Dagnino, B., Gariel-Mathis, M.-A., Poort, J., van der Togt, C., & Roelfsema, P. R. (2014). Alpha and gamma oscillations characterize feedback and feedforward processing in monkey visual cortex. *Proceedings of the National Academy of Sciences*, *111*(40), 14332–14341. https://doi.org/10.1073/pnas.1402773111

van Kerkoerle, T., Self, M. W., & Roelfsema, P. R. (2017). Layer-specificity in the effects of attention and working memory on activity in primary visual cortex. *Nature Communications*, *8*(1), 13804. https://doi.org/10.1038/ncomms13804

VanRullen, R. (2016). Perceptual Cycles. *Trends in Cognitive Sciences*, *20*(10), 723–735. https://doi.org/10.1016/j.tics.2016.07.006

VanRullen, R. (2018, March 23). Perceptual Rhythms. In J. T. Wixted (Ed.), *Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience* (pp. 1–44). John Wiley & Sons, Inc. https://doi.org/10.1002/9781119170174.epcn212

Varela, F. J., Toro, A., Roy John, E., & Schwartz, E. L. (1981). Perceptual framing and cortical alpha rhythm. *Neuropsychologia*, *19*(5), 675–686. https://doi.org/10.1016/0028-3932(81)90005-1

Venables, P. H. (1960). Periodicity in Reaction Time. *British Journal of Psychology*, *51*(1), 37–43. https://doi.org/10.1111/j.2044-8295.1960.tb00722.x

Vinje, W. E., & Gallant, J. L. (2000). Sparse Coding and Decorrelation in Primary Visual Cortex During Natural Vision. *Science*, *287*(5456), 1273–1276. https://doi.org/10.1126/science.287.5456.1273

Vogels, R. (1990). Population coding of stimulus orientation by striate cortical cells. *Biological Cybernetics*, *64*(1), 25–31. https://doi.org/10.1007/BF00203627

Wagenmakers, E.-J. (2007). A practical solution to the pervasive problems of p values. *Psychonomic Bulletin & Review*, *14*(5), 779–804. https://doi.org/10.3758/BF03194105

Wagenmakers, E.-J., & Farrell, S. (2004). AIC model selection using Akaike weights. *Psychonomic Bulletin & Review*, *11*(1), 192–196. https://doi.org/10.3758/BF03206482

Walsh, E. G. (1952). Visual reaction time and the α-rhythm, an investigation of a scanning hypothesis. *The Journal of Physiology*, *118*(4), 500–508. https://doi.org/10.1113/jphysiol.1952.sp004811

Walter, V. J., & Walter, W. G. (1949). The central effects of rhythmic sensory stimulation. *Electroencephalography and Clinical Neurophysiology*, *1*(1), 57–86. https://doi.org/10.1016/0013-4694(49)90164-9

Watson, A. B. (2017). QUEST+: A general multidimensional Bayesian adaptive psychometric method. *Journal of Vision*, *17*(3), 10. https://doi.org/10.1167/17.3.10

Wei, X.-X., & Stocker, A. A. (2015). A Bayesian observer model constrained by efficient coding can explain 'anti-Bayesian' percepts. *Nature Neuroscience*, *18*(10), 1509–1517. https://doi.org/10.1038/nn.4105

Wei, X.-X., & Stocker, A. A. (2016). Mutual Information, Fisher Information, and Efficient Coding. *Neural Computation*, *28*(2), 305–326. https://doi.org/10.1162/NECO_a_00804

Wei, X.-X., & Stocker, A. A. (2017). Lawful relation between perceptual bias and discriminability. *Proceedings of the National Academy of Sciences*, *114*(38), 10244–10249. https://doi.org/10.1073/pnas.1619153114

Williams, J. G., Harrison, W. J., Beale, H. A., Mattingley, J. B., & Harris, A. M. (2024). Effects of neural oscillation power and phase on discrimination performance in a visual tilt illusion. *Current Biology*. https://doi.org/10.1016/j.cub.2024.03.014

Wilson, H. R., & Cowan, J. D. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal*, *12*(1), 1–24. https://doi.org/10.1016/S0006-3495(72)86068-5

Wilson, M. A., Sumera, A., Taylor, L. W., Meftah, S., McGeachan, R. I., Modebadze, T., Jayasekera, B. A. P., Cowie, C. J. A., LeBeau, F. E. N., Liaquat, I., Durrant, C. S., Brennan, P. M., & Booker, S. A. (2025). Phylogenetic divergence of GABAB receptor signaling in neocortical networks over adult life. *Nature Communications*, *16*(1), 4194. https://doi.org/10.1038/s41467-025-59262-8

Womelsdorf, T., Valiante, T. A., Sahin, N. T., Miller, K. J., & Tiesinga, P. (2014). Dynamic circuit motifs underlying rhythmic gain control, gating and integration. *Nature Neuroscience*, *17*(8), 1031–1039. https://doi.org/10.1038/nn.3764

Wutz, A., Melcher, D., & Samaha, J. (2018). Frequency modulation of neural oscillations according to visual task demands. *Proceedings of the National Academy of Sciences*, *115*(6), 1346–1351. https://doi.org/10.1073/pnas.1713318115

Xing, J., & Heeger, D. J. (2001). Measurement and modeling of center-surround suppression and enhancement. *Vision Research*, *41*(5), 571–583. https://doi.org/10.1016/S0042-6989(00)00270-4

Yuille, A., & Bülthoff, H. (1996). Bayesian decision theory and psychophysics. In D. C. Knill & W. Richards (Eds.), *Perception as Bayesian Inference* (pp. 123–162). Cambridge University Press. https://doi.org/10.1017/CBO9780511984037.006

Zazio, A., Ruhnau, P., Weisz, N., & Wutz, A. (2021). Pre-stimulus alpha-band power and phase fluctuations originate from different neural sources and

exert distinct impact on stimulus-evoked responses. *European Journal of Neuroscience*, *55*(11–12). https://doi.org/10.1111/ejn.15138

Zhang, H., Morrone, M. C., & Alais, D. (2019). Behavioural oscillations in visual orientation discrimination reveal distinct modulation rates for both sensitivity and response bias. *Scientific Reports*, *9*(1), 1115. https://doi.org/10.1038/s41598-018-37918-4

Zhang, L.-Q., Mao, J., Aguirre, G. K., & Stocker, A. A. (2025). The tilt illusion arises from an efficient reallocation of neural coding resources at the contextual boundary. *Proceedings of the National Academy of Sciences*, *122*(17), e2421565122. https://doi.org/10.1073/pnas.2421565122

Zhigalov, A., & Jensen, O. (2020). Alpha oscillations do not implement gain control in early visual cortex but rather gating in parieto-occipital regions. *Human Brain Mapping*, *41*(18), 5176–5186. https://doi.org/10.1002/hbm.25183

Zhou, J., Duong, L. R., & Simoncelli, E. P. (2024). A unified framework for perceived magnitude and discriminability of sensory stimuli. *Proceedings of the National Academy of Sciences*, *121*(25), e2312293121. https://doi.org/10.1073/pnas.2312293121

Zhou, Y. J., Iemi, L., Schoffelen, J.-M., Lange, F. P. de, & Haegens, S. (2021). Alpha oscillations shape sensory representation and perceptual sensitivity. *Journal of Neuroscience*, *41*(46), 9581–9592. https://doi.org/10.1523/JNEUROSCI.1114-21.2021

Appendix A
Ethics Approval Letter

THE UNIVERSITY OF QUEENSLAND
AUSTRALIA
CREATE CHANGE
Research Ethics and Integrity

Human Research Ethics Approval

Project Number:	2021/HE002284
Project Title:	Understanding how neural oscillatory phase affects perception and attention
Version:	8.01
Chief Investigator:	Dr Anthony Harris
	Queensland Brain Institute
Co-Investigator(s)	Dr Amanda Robinson Mr Ben Carman Mr Henry Beale Professor Jason Mattingley Ms Jessica Williams Ms Jasmin Patel Ms Jessica Elliott Miss Kali Chidley Ms Leah Koger Ms Morgan McIntyre Mr Matthew O'Donohue Mr Samuel Campbell
Funding Body (UQ ref#):	
Approving Committee:	University of Queensland Human Research Ethics Committee A
Approval End Date:	31 Dec 2026
Date of Approval:	Friday, 29 May, 2026

University of Queensland Human Research Ethics Committee A confirms that this project meets the requirements of the National Statement on Ethical Conduct in Human Research (2023). The University's human research ethics committees are organised and operate in accordance with the National Statement on Ethical Conduct in Human Research (2023).

Approved Documents

Document Type	File Name	Document Tile	Application Version	Document Version	Last Modified
Project Protocol	2021_HE002284_v12 - Project_Description.docx	2021_HE002284_v12 - Project_Description.docx	8.01	4	28/05/2026 12:04:36 PM
Change Tracking	2021_HE002284 v7_02 - v8_01 Changes.pdf	2021/HE002284 v7_02 - v8_01 Changes	8.01	1	28/05/2026 12:04:42 PM

Application Attachment	2021_HE002284_v12 - Consent form.docx	Consent form	8.01	5	28/05/2026 12:04:37 PM
Application Attachment	2021_HE002284_v12 - Consent form_tracked.docx	Consent form, tracked changes	8.01	1	28/05/2026 12:04:38 PM
Application Attachment	flyer_v2.pdf	Flyer for advertising study	7.02	2	28/05/2026 12:04:38 PM
Application Attachment	2021_HE002284_v12 - MRI Safety form.pdf	MRI Safety Form - no changes	8.01	5	28/05/2026 12:04:38 PM
Application	Output Form.pdf	Output Form	8.01	14	28/05/2026 12:04:36 PM
Application Attachment	2021_HE002284_v12 - PLS - S1.docx	Plain Language Statement - Study 1	8.01	5	28/05/2026 12:04:37 PM
Application Attachment	2021_HE002284_v12 - PLS - S1_tracked.docx	Plain Language Statement - Study 1, tracked changes	8.01	1	28/05/2026 12:04:38 PM
Application Attachment	2021_HE002284_v12 - PLS - S2.docx	Plain Language Statement - Study 2	8.01	5	28/05/2026 12:04:37 PM
Application Attachment	2021_HE002284_v12 - PLS - S2_tracked.docx	Plain Language Statement - Study 2, tracked changes	8.01	1	28/05/2026 12:04:38 PM
Application Attachment	2021_HE002284_v12 - PLS - S3.docx	Plain Language Statement - Study 3	8.01	5	28/05/2026 12:04:37 PM
Application Attachment	2021_HE002284_v12 - PLS - S3_tracked.docx	Plain Language Statement - Study 3, tracked changes	8.01	1	28/05/2026 12:04:38 PM
Application Attachment	2021_HE002284_v12 - Project_Description.docx	Project description	8.01	5	28/05/2026 12:04:37 PM

THE UNIVERSITY OF QUEENSLAND
AUSTRALIA
CREATE CHANGE
Research Ethics and Integrity

Application Attachment	2021_HE002284_v12 - Project_Description_track.docx	Project description, tracked changes	8.01	1	28/05/2026 12:04:38 PM

Ethics committee

University of Queensland Human Research Ethics Committee A
EC00456
The University of Queensland

Additional Notes to HREC Approval

1. The University of Queensland HREC is constituted and operates in accordance with the National Health and Medical Research Council's *National Statement on Ethical Conduct in Human Research (2023)*, NHMRC and Universities Australia *Australian Code for the Responsible Conduct of Research (2018)* and the *CPMP/ICH Note for Guidance on Good Clinical Practice*.

2. If a Public Health Act (PHA) application is applicable, please visit the Health and Medical Research Unit website at:
   http://www.health.qld.gov.au/ohmr/html/regu/aces_conf_hth_info.asp
   Researchers are reminded that they require a current ethics application to utilise data provided through a PHA application.

3. In accordance with Section 5.5.6 (b) of the National Statement, the Chief Investigator will report to the HREC annually for the duration of the project. A final report is to be submitted on completion of the study.

4. The Chief Investigator will immediately report anything that might warrant review of ethical approval of the project, in the specified format, including:
   - Unforeseen events that might affect continued ethical acceptability of the project.
   - Serious Adverse Events that materially impact on the continued ethical acceptability of the project.

   Any other incident attributable to the research affecting the welfare and/or health of participants or researchers should be promptly reported through UQSafe and the relevant reviewing HREC.

5. Amendments to any part of the approved protocol (including change of Investigator/s), documents, or questionnaires attached to the clearance must be submitted to the HREC and approval granted before the changes are implemented.

6. All clinical trials need to be registered on a World Health Organization (WHO) approved clinical trials registry (for example http://www.anzctr.org.au).

7. The Chief Investigator must determine whether the study needs to be declared/notified to UQ Insurance Services as per Fact Sheet on Insurance Services website (refer flow chart at Appendix 1).[1]

8. The Chief Investigator is responsible and accountable for full compliance of the protocol by all investigators including the collection, use, storage and disclosure of data as required by UQ policies and procedures.

9. The HREC reserves the right to visit the research site and view materials at any time, and to conduct a full audit of the project.

[1] Despite international and disciplinary norms, a wide range of studies involving humans can be considered a "Clinical Trial" for insurance purposes at UQ. E.g., in some cases, this can include epidemiological surveys, population level studies, blood sampling studies (e.g., seroprevalence), psychophysiological outcome studies, and parenting intervention evaluations.

THE UNIVERSITY OF QUEENSLAND
AUSTRALIA
CREATE CHANGE
Research Ethics and Integrity

Additional Notes to LNR Panel Approval

1. Research reviewed and approved by LNR Panel meets National Statement Requirements as per s5.1.18 – s5.1.21

2. In accordance with Section 5.5.6 (b) of the National Statement, the Chief Investigator will report to the LNR Panel annually for the duration of the project. A final report is to be submitted on completion of the study.

3. The Chief Investigator will immediately report anything which might warrant review of ethical approval of the project, in the specified format, including:
   - Unforeseen events that might affect continued ethical acceptability of the project.
   - Serious Adverse Events that materially impact on the continued ethical acceptability of the project.

   Any other incident attributable to the research affecting the welfare and/or health of participants or researchers should be promptly reported through UQSafe and the relevant LNR Panel.

4. Amendments to any part of the approved project description (including change of Investigator/s), documents, or questionnaires attached to the application must be submitted to the HREC and approval granted before the changes are implemented.

5. The Chief Investigator is responsible and accountable for full compliance of the protocol by all investigators including the collection, use, storage and disclosure of data as required by UQ policies and procedures.

6. The LNR Panel reserves the right to visit the research site and view materials at any time, and to conduct a full audit of the project.