Effect of stimulus repetition, color and context on visual cortical gamma

Der Effekt von Stimuluswiederholung, Farbe und Kontext auf visuelle kortikale Gamma-Oszillationen

Dissertation

zur Erlangung des Doktorgrades
der Naturwissenschaften

vorgelegt beim Fachbereich Biowissenschaften
der Johann Wolfgang Goethe-Universität
in Frankfurt am Main von

Alina Peter
aus
Dortmund

Frankfurt (2019) D(30)

GOETHE UNIVERSITÄT FRANKFURT AM MAIN

Ernst Strüngmann Institute for Neuroscience in Cooperation with Max Planck Society

Vom Fachbereich 15 der Johann Wolfgang von Goethe-Universität als Dissertation angenommen.
Dekan: Prof. Dr. Sven Klimpel
Gutachter: Prof. Dr. Manfred Kössl und Prof. Dr. Pascal Fries
Datum der Disputation: 27.02.2020

Contents

Abstract

Humans and other primates are highly visual animals. Our daily visual activities such as recognizing familiar faces, interacting with objects, or reading, are supported by an extensive system of interacting brain areas. The interactions between the many individual nerve cells both within and between brain areas need to be coordinated. One possible solution to achieve flexible coordination between cells in the network is rhythmic activity, or oscillations. The focus of the thesis will be activity in the largest visual area, V1, in non-human primates. In V1, high-frequency activity, so-called gamma-band activity ("gamma", ca. 30-90 Hz) can be frequently observed and has been suggested to play a role in coordinating activity in the visual system. In Chapter 1, the coordination problem, the primate visual system and gamma-band oscillations are introduced in detail. The following chapters explore the dependence of gamma on contextual influences. Does V1 use contextual information to optimize coordination? In the first part, the short-term consequences of repeated encounters with visual stimuli on V1 responses are explored (Chapters *2* and *3*). Inspired by results from colored, naturalistic images in the first part, the second part tests the dependence of gamma on spatial and chromatic stimulus aspects (Chapters *4* and *5*).

Stimulus repetition is a simple yet powerful way to tap into our brains' ability to learn and adapt to our environment. Repeated presentation of a visual stimulus tends to decrease responses to this stimulus. Is this accompanied by changes in the coordination of brain activity? In Chapter *2*, the stimulus-specificity of repetition effects on gamma was tested using naturalistic stimuli. V1 is most typically studied using black-and-white, artificial stimuli that are very familiar to the animals. Here, colored natural images were repeatedly presented that were initially novel to the animals, to provide a wider and more naturalistic range of stimulation. Both multi-unit spiking activity (MUA) and gamma showed stimulus-specific repetition effects. MUA responses decreased most strongly for initial repetitions and less for later repetitions. In contrast, gamma could increase or decrease for initial repetitions, but tended to increase for later repetitions. This points to the operation of multiple plasticity mechanisms. One process may rapidly decrease MUA and gamma and be related to initial novelty or adaptation. The other increases gamma, is active for more repetitions, and could constitute a form of refinement of coordination over time. Morevoer, based on the spacing of stimulus repetitions, stimulus memory in V1 persisted for tens of seconds.

In the following Chapter *3*, the stimulus location specificity and persistence of the repetition effects for longer timescales were tested. To this end, the observation that the increase in gamma with repetition was strongest for the first tens of repetitions was used to test for location specificity and memory. Using simple artificial stimuli that were repeated many times at two alternating locations, both location specificity and memory on the order of minutes was observed. Due to the structure of the primate visual system, location specificity suggests that the repetition effects involve early to mid-level visual areas such as V1. Memory for previous stimulus presentations on

the order of minutes has not been previously reported for V1 gamma. Taken together, these experiments demonstrate short-term plasticity of gamma that is stimulus- and location specific and persists on the timescale of minutes.

In Chapter *2*, the average gamma-band response to the large, naturalistic stimuli was highly stimulus dependent. Relative increases in gamma-band activity scaled between tens and thousands of percent change depending on the stimulus. Particularly the color of the stimuli appeared to play a strong role, although the stimulus set was too limited and uncontrolled to draw strong conclusions. In Chapters *4* and *5*, underlying mechanisms for the stimulus specificity of gamma were explored using more well-controlled, artificial stimuli that varied in color and spatial structure.

Much of vision relies on the analysis of spatial structure. Each nerve cell in V1 only responds to visual stimuli in a particular, small part of the visual field, its so-called "receptive field" (RF). Compared to isolated RF stimulation, nearby cells that are stimulated by a similar structure from different parts of visual space can show response decreases, commonly known as "surround suppression", and may show coordinated activity in the gamma band. In Chapter *3*, responses to large, uniformly colored disks are contrasted with responses to black or white (achromatic) disks. A first experiment showed that gamma-band responses were stronger for colored than achromatic stimuli, whereas MUA responses could decrease below baseline for colored stimuli. To test whether these phenomena were related to surround suppression, stimulus size was manipulated in a second experiment. When stimuli were of sufficient size to induce surround suppression, clear gamma-band responses emerged. Surround suppression and gamma were stronger for chromatic stimuli. However, the change of stimulus size could have changed not only surround suppression but also stimulus saliency. Therefore, in a third experiment, the overall size of the stimulus was kept constant, and the spatial structure of the stimulus was manipulated. In comparison to uniform, predictable stimulus structure, mismatches between the center of the stimulus and the surrounding visual space led to strong increases in MUA responses and strong decreases in gamma-band activity. These effects were restricted to the recording sites with RFs at the mismatch location. These experiments underpin the strong role of both spatial structure and color for gamma in V1.

In Chapter *4*, responses to different color hues are studied in more detail. Gamma response strength depended on hue, being strongest for red compared to blue and green stimuli when measured with a gray background. To better understand the underlying mechanisms of the differential responses, the spatio-temporal context in the form of the background color was manipulated. Background color had a strong influence on gamma strength. Using differently colored backgrounds, different parts of the color signaling pathways could be adapted. Response differences to different color hues could be explained well with a model that incorporates differences in adaptation between pathways involving long- compared to medium-wavelength cone signals.

Taken together, these experiments indicate a strong role of both spatial context (stimulus size and structure) and temporal context and drive (repetition, adaptation) for the generation of gamma-band activity in V1. Functional implications of these dependencies are considered in the final Chapter *6*, and a role for gamma-band synchronization in a coding regime for visual inputs that generate strong drive and high predictability is suggested.

Keywords:
gamma-band oscillations, stimulus repetition, surround suppression, color, V1

Ausführliche deutsche Zusammenfassung

Menschen und andere Primaten verfügen über hervorragende visuelle Wahrnehmungsfähigkeiten. Diese basieren auf einem komplexen System interagierender Hirnareale, welches uns zum Beispiel ermöglicht, Gesichter zu erkennen, mit Gegenständen umzugehen oder zu lesen. Aufgrund der vielen Verbindungen zwischen Nervenzellen sowohl innerhalb eines Hirnareals als auch zwischen Arealen interagieren Zellen stark miteinander, sodass die Aktivität verschiedener Zellen zeitlich koordiniert werden muss. Eine mögliche Grundlage flexibler Koordination in einem Netzwerk von Zellen ist ihre Synchronisierung mittels rhythmischer Aktivität (Oszillationen).

Diese Arbeit beschäftigt sich mit der Aktivität im größten visuellen Hirnareal, genannt V1, bei nicht-menschlichen Primaten (Rhesusaffen). In V1 ist hochfrequente Aktivität im sogenannten Gamma-Band (ca. 30-90 Hz, „Gamma“) ein häufiges Merkmal von neuronalen Antworten auf visuelle Stimuli. Da synchronisierte Aktivität im Gamma-Band spezifisch zwischen Zellen auftreten kann, die den selben visuellen Stimulus verarbeiten, wurde Gamma eine funktionelle Rolle in der Koordination von Hirnaktivität zugesprochen. In Kapitel *1* werden das visuelle System des Primaten, das Koordinationsproblem zwischen Zellen und Gamma-Band Oszillationen ausführlich eingeführt.

Diese Dissertation behandelt die Abhängigkeit von Gamma in V1 von verschiedenen kontextuellen Einflüssen. Nutzt V1 kontextuelle Informationen, um seine Koordination zu optimieren? Im ersten Teil wird untersucht, inwiefern Stimuluswiederholungen Veränderungen in den Antworten in V1 und insbesondere in Gamma hervorrufen (Kapitel *2* und *3*). Inspiriert von den starken Antworten im Gamma-Band auf farbige, natürliche Stimuli, wird im zweiten Teil der Zusammenhang von Gamma zu bestimmten Stimuluseigenschaften getestet, insbesondere ihrer räumlichen Struktur und Farbe (Kapitel *4* und *5*). Die Zusammenhänge zwischen räumlichen und zeitlichen Einflüssen auf Gamma werden im letzten Kapitel *6* diskutiert und in den breiteren Zusammenhang mit dem aktuellen Forschungsstand und vorherrschenden Theorien zu visueller Kodierung gesetzt.

Einfluss von Stimuluswiederholung auf visuelle Gamma-Band Aktivität

Unser Gehirn befindet sich in einem ständigen Prozess des Lernens und der Anpassung an das momentane Umfeld. Zum Beispiel reagiert das Gehirn auf die wiederholte Präsentation eines Stimulus nicht identisch, sondern vermindert typischerweise seine Antworten in Form von Feuerraten. Erfolgt gleichzeitig eine verstärkte Koordination der verbleibenden Antworten? Eine frühere Studie des Labors fand starke Hinweise auf solche verbesserte Koordination spezifisch im Gamma-Band sowohl innerhalb als auch zwischen Hirnarealen im Verlauf von hunderten Präsentationen von künstlichen

Stimuli. Jedoch blieb unklar, inwiefern diese Verbesserung der Koordination spezifisch für einen bestimmten Stimulus war. In Kapitel *2* wird die Stimulusspezifität von Wiederholungseffekten auf Gamma mittels natürlicher Stimuli getestet. Antworten in V1 werden meist mittels künstlicher Schwarzweißstimuli untersucht, welche die Tiere schon oftmals zuvor gesehen haben. In dieser Studie wurden farbige Fotografien von Objekten verwendet, welche die Tiere zu Beginn des Experiments das erste Mal sahen. Mit Hilfe dieser Stimuli kann getestet werden, ob der ursprünglich beobachtete Wiederholungseffekt unter breiteren und natürlicheren Bedingungen auftritt.

Der Effekt von Stimuluswiederholung zeigte Stimulusspezifität, sowohl für die Feuerraten von Zellgruppen (Multi-Unit-Aktivität, MUA) als auch für Gamma-Band Aktivität. Dies zeigte sich sowohl durch eine Korrelationsanalyse von Wiederholungseffekten über die Ableitetage hinweg als auch durch lineare Regressionsmodelle. Im Gamma-Band zeigte sich die Stimulusspezifität der Wiederholungseffekte sowohl für die Stärke der Antworten im lokalen Feldpotential (LFP, d.h. vor allem synchronisierte extrazelluläre Signale wie synaptische Aktivität), als auch in der Koordination der MUA zum LFP. Niedrigere Frequenzen zeigten keine Stimulusspezifität.

Stimulus-spezifische Wiederholungseffekte in der MUA traten früh (< 100 ms nach Beginn der Stimulation) auf und hielten für mehrere Hundert Millisekunden an. Die geringe Latenz der Effekte spricht für eine Rolle von V1 oder seinen Inputs, und gegen eine Rolle von höheren Arealen oder kognitiven Effekten wie Aufmerksamkeit als Erklärung für den Wiederholungseffekt. Des Weiteren basierten die Wiederholungseffekte nicht allein auf direkten Wiederholungen, sondern zeigten sich auch über Unterbrechungen durch mehrere andere Stimuli hinweg.

MUA wurde generell durch Stimuluswiederholung vermindert. Dies galt insbesondere für die ersten Wiederholungen. Dabei verringerten sich die Antworten in der MUA für die Stimuli und Ableitestellen am meisten, welche anfänglich stärker waren. Dies spricht gegen eine verbesserte Prozessierung durch selektiven Wegfall schwach antwortender Zellen. Die Veränderung von Gamma-Band Aktivität durch Stimuluswiederholung zeigte dagegen ein komplexeres Bild. Für die ersten Wiederholungen konnte die Stärke von Gamma-Band Antworten in Abhängigkeit vom Stimulus ansteigen oder abfallen. Bei späteren Wiederholungen dominierten Anstiege die Antworten im Gamma-Band. Stimuli, die schon während der ersten Präsentationen Anstiege zeigten, zeigten in der Tendenz auch für spätere Präsentationen stärkere Anstiege. Dieses Muster an Effekten weist auf zwei verschiedene unterliegende Plastizitätsprozesse hin, die in Abhängigkeit vom Stimulus unterschiedlich stark wirken. Der eine Prozess, der ähnlich wie bei der MUA am stärksten die anfänglichen Wiederholungen betrifft, führt zu verminderten Gamma-Band Antworten. Starke Feuerraten und Gamma-Band Antworten für die erste Präsentation könnten zum Beispiel die Neuheit eines Stimulus reflektieren. Der zweite Prozess verstärkt Gamma-Band Antworten, nicht nur für die ersten Wiederholungen, und könnte einer durch stärkere Koordination verbesserten Prozessierung des Stimulus dienen. Es ist bemerkenswert, dass die alleinige Analyse von Feuerraten wenig Hinweise auf anhaltende Plastizitätsprozesse nach den ersten Wiederholungen lieferte, diese aber in der Veränderung der Koordination der Netzwerkaktivität deutlich wurden.

Die Wiederholungseffekte von Gamma-Band Antworten hingen ebenfalls von der anfänglichen Stärke der Antworten auf einen Stimulus ab. Wogegen die MUA für einen Stimulus stärker abfiel, wenn dieser anfänglich starke Antworten induzierte, zeigte sich das gegenteilige Bild für Gamma-Band Antworten: je stärker die anfänglichen Antworten, desto stärker war auch der Anstieg der Antworten mit Stimuluswiederholung. Die Stärke der Veränderung der Koordination der V1-Antworten hing

demnach vom Stimulus ab. Ein verbessertes Verständnis der generellen Stimulusabhängigkeit von Gamma-Band Antworten kann daher helfen zu verstehen, welche Stimuli mit Wiederholung stärkere Plastizität in Gamma zeigen (siehe Kapitel *4* und 5). Da zwischen der Präsentation eines Stimulus und seiner Wiederholung ca. 2-20 Sekunden vergehen konnten (und unterschiedliche Stimuli gezeigt wurden), zeigt dieses Experiment zudem, dass Information in V1 über mindestens solche Zeiträume integriert werden kann. In Kapitel *3* wurde u.a. die Dauer des Effektes weiter untersucht.

In Kapitel *3* wurde zunächst getestet, inwiefern der Anstieg von Gamma-Band Antworten durch viele Wiederholungen für einen bestimmten Ort im visuellen Feld spezifisch ist. Das Experiment wurde gleichzeitig genutzt, um zu testen, ob V1 für längere Zeiträume anhaltende Veränderungen aufgrund von Stimuluswiederholung zeigt. Dazu wurde die Tatsache genutzt, dass die Rate des Anstiegs von Gamma-Band Antworten sich bei sehr vielen Wiederholungen (>> 10) abschwächt. Durch die vielmalig wiederholte Präsentation einfacher künstlicher Stimuli an zwei alternierenden Orten wurde die räumliche Spezifität festgestellt, indem die Rate des Anstiegs bei vorheriger Stimulation am gleichen mit bei vorheriger Stimulation an einem anderen Ort verglichen wurde. Zudem zeigte sich, dass der Anstieg in Gamma-Band Antworten bei einer Unterbrechung von mehreren Minuten erhalten blieb – Gamma reduzierte sich nicht zurück auf das Niveau zu Beginn der Stimuluswiederholungen.

Da im visuellen System nur frühe Hirnareale räumlich spezifische Antworten geben, ist die räumliche Spezifität des Effekts ein Hinweis darauf, dass der Effekt aus V1 selbst oder benachbarten Arealen stammt. Eine anhaltende Erhöhung von Gamma trotz minutenlanger Unterbrechung ist ein bisher unbekannter Effekt, der darauf hinweist, dass länger anhaltende Plastizitätsmechanismen involviert sind.

Zusammengenommen zeigen die Experimente in Kapitel *2* und 3 einen starken und über mittlere Zeiträume anhaltenden Effekt von Stimuluswiederholung auf V1 Antworten, sowohl in der MUA als auch in der Gamma-Band Koordinierung. Die unterschiedlichen Effekte von anfänglichen im Vergleich zu späteren Wiederholungen auf Gamma deuten auf das simultane Wirken mehrerer Plastizitätsprozesse hin. Diese könnten beispielsweise der schnellen Prozessierung neuer, und der sich langsam verbessernden Prozessierung sich stetig wiederholender Stimuli dienen.

Einfluss von räumlicher Struktur, Farbe und Kontext auf visuelle Gamma-Band Aktivität

In Kapitel *2* fiel auf, dass die Stärke der Antworten im Gamma-Band auf natürliche Stimuli stark vom individuellen Stimulus abhing. Der relative Anstieg von Gamma durch Stimuli variierte zwischen ein paar wenigen bis ein paar Tausend Prozent. Insbesondere die Farbe der Stimuli schien dabei eine starke Rolle zu spielen. Allerdings waren die Stimuli fuer eine eindeutige Untersuchung unzureichend. Dies lag insbesondere an der unkontrollierten Kovarianz verschiedener Stimulusaspekte bei natürlichen Stimuli. In Kapitel *4* und *5* wurden daher die Mechanismen, die Gamma stimulusabhängig machen, näher untersucht. Dafür wurden stärker kontrollierte, künstliche Stimuli verwendet, deren Farbe und räumliche Struktur systematisch verändert wurde.

Visuelle Wahrnehmung beruht auf der erfolgreichen Erfassung von räumlichen Beziehungen und Strukturen zwischen verschiedenen Teilen der visuellen Welt. Jede Nervenzelle in V1 ist auf einen bestimmten, kleinen Teil des visuellen Felds spezialisiert, ihr sogenanntes „rezeptives Feld“ (RF). Im Vergleich zur alleinigen Stimulati-

on des RF sind die Feuerraten von Nervenzellen geringer, wenn benachbarte Zellen gleichzeitig durch einen zusammenhängenden oder ähnlichen Stimulus angeregt werden. Dieses Phänomen wird “Surround Suppression” (Unterdrückung von Antworten durch das nahe visuelle Umfeld) genannt und wurde für Schwarzweißstimuli bereits durch vorherige Forschung in Zusammenhang mit Gamma gestellt.

In Kapitel 3 wurden Antworten in V1 auf große, einheitlich farbige Stimuli mit Antworten auf schwarze oder weiße (achromatische) Stimuli verglichen. In einem ersten Experiment zeigte sich, dass Gamma-Band Aktivität für farbige Stimuli deutlich stärker war, wogegen die MUA noch unter Antworten auf einen grauen Hintergrund abfiel. Diese geringe MUA könnte ein Hinweis auf eine Rolle von Surround Suppression sein. Um eine Verbindung zu Surround Suppression zu testen, wurde daher in einem zweiten Experiment die Größe der Stimuli systematisch verändert. Sobald Stimuli groß genug waren um Surround Suppression in der MUA zu bewirken, wurden auch Antworten im Gamma-Band deutlich. Sowohl Surround Suppression als auch Gamma waren für farbige Stimuli stärker. Die Veränderung der Größe der Stimuli brachte allerdings nicht nur Veränderungen in Surround Suppression mit sich, sondern aller Wahrscheinlichkeit auch Veränderungen in der Salienz und damit Aufmerksamkeit für die Stimuli. In einem dritten Experiment wurde daher statt der Größe die innere Struktur der Stimuli verändert. Im Vergleich zur Stimulation mit einheitlicher und daher räumlich vorhersagbarer Farbe führte eine uneinheitliche Stimulation von einem RF und seinem Umfeld zu starken Anstiegen in der MUA, und gleichzeitig starken Abfällen von Gamma. Diese Effekte waren sehr spezifisch für Antworten auf den Ort einer Unstimmigkeit zwischen RF und dem Umfeld des RF. Die Experimente zeigen insgesamt eine starke Rolle sowohl von der Vorhersagbarkeit oder Einheitlichkeit der räumlichen Struktur als auch der Farbe auf V1 Aktivität. Die starken Antworten auf farbige im Vergleich zu achromatischen Stimuli werden aufgrund der bekannten Physiologie des Farbsystems und der Mechanismen von Gamma für andere Stimuli als eine Folge von stärkerem Input (Anregung) ausgelegt, unabhängig vom spezifischen Farbton.

In Kapitel *4* wurden die unterschiedlichen Antworten auf *verschiedene* Farben genauer untersucht. Antworten auf farbige Stimuli über den gesamten verfügbaren Farbraum des Monitors, und mit verschiedenen Helligkeiten, waren stärker als auf achromatische Stimuli. Verschiedene Kontrollexperimente zeigten, dass dies nicht durch unterschiendlichen Luminanz- oder Farbkontrast bedingt war. Die Stärke von Gamma war allerdings vom genauen Farbton abhängig. Auf grauem Hintergrund waren die Antworten auf rote und blaue Töne am stärksten. Um die Ursprünge dieser Unterschiede besser zu verstehen, wurde das Experiment mit verschiedenen Hintergrundfarben wiederholt. Diese Veränderung des Kontexts hatte große Auswirkungen auf die Stärke der Antworten im Gamma-Band. Die Veränderung des Hintergrunds hatte den experimentellen Zweck, verschiedene Kanäle der Farbverarbeitung gezielt zu adaptieren. Ein einfaches Modell konnte die Veränderungen von Gamma-Band Antworten mit dem Hintergrund durch eine unterschiedlich starke Adaptation der Signale von mittleren im Vergleich zu langen Wellenlängen erfassen.

Kapitel *2* bis *4* zeigen, dass Gamma stark vom Kontext sowohl räumlicher als auch zeitlicher Natur abhängt. Die Zusammenhänge zwischen räumlichen und zeitlichen kontextuellen Einflüssen auf Gamma werden im letzten Kapitel *6* diskutiert.

Fazit

Insgesamt demonstrieren diese Experimente, dass Kontext sowohl räumlicher (Größe und Struktur der Stimuli) als auch zeitlicher (Stimuluswiederholung, Adaptation) Form eine große Rolle für Aktivität in V1 spielt, insbesondere für Gamma und damit für die Koordination von Aktivität in V1. Im breiteren Zusammenhang mit dem aktuellen Forschungsstand zeigen die vorliegenden Ergebnisse eine besondere Rolle von sowohl räumlicher als auch zeitlicher Vorhersagbarkeit für Gamma. Dies setzt Gamma-Band Aktivität in direkten Zusammenhang zu aktuellen Theorien visueller Kodierung, die eine effiziente Kodierung vorhersagbarer Information postulieren.

Zusätzlich deuten die Effekte von Farbe, bekannte Effekte von Stimuluskontrast, und eine Abhängigkeit der Widerholungseffekte von der ursprünglichen Stärke der Gamma-Band Aktivität, auf eine Rolle der Stärke der Anregung von V1 durch visuelle Stimuli auf die Stärke der Gamma-Band Aktivität hin. Aufgrund der Abhängigkeit von der Stärke der Anregung durch den Stimulus und vom Kontext lässt sich Gamma-Band Synchronisierung als eine bestimmte Form von Kodierung visueller Stimuli interpretieren, die insbesondere bei starken, vorhersagbaren Stimuli auftritt. Diese koordinierte Form von Kodierung könnte sowohl besonders effizient als auch akkurat sein, und somit besondere Auswirkungen auf die weitere Verarbeitung in späteren Hirnarealen und letztlich auf die Wahrnehmung haben.

Stichworte:
Gamma-Band Oszillationen, Stimuluswiederholung, Surround Suppression, Farbe, V1

Chapter 1

General Introduction

1.1 Acknowlegdements for Chapter 1

For this chapter, I produced some figures to illustrate basic properties of the laminar structure of V1. For this, I used data collected by Eric Lowet and Mark Roberts in the lab of Peter de Weerd. In particular, as part of this project I developed an algorithm based on parallel tempering, a variant of Monte Carlo methods, to align data collected with a laminar probe between days, and integrated this alignment with a receptive-field based assignment of visual areas. The method is described in detail in the resulting publication, Lowet, Roberts, *Peter*, Gips and De Weerd, eLife (2017). All analyses and figures based on the collected data presented in this chapter are produced by myself.

1.2 Vision as a complex coordination problem

"We're least aware of what our minds do best" Minsky (1986), p.29

Much of our conscious experience, our memories and our interaction with the world is shaped by vision. Imagine that you want to take a bottle that a friend offers you on the other side of the room. This involves solving visual tasks like the recognition and manipulation of objects (the bottle), navigation in space (find your way around the room, avoiding obstacles), or judging a face as familiar (who is giving you the bottle). Such tasks appear simple to us. Evidence that visual tasks are in fact not trivial originally came from patients with particular brain lesions who show deficits in highly specific tasks (Mishkin et al., 1983; Goodale and Milner, 1992). It is also evident from the immense difficulty encountered when trying to teach machines how to see: Computers could play chess long before they could tell apart a picture of a cat from that of a turtle. Vision is therefore a difficult task that humans solve with surprising ease. The discrepancy between perceived and true difficulty of the task may be resolved when considering the sophistication of the underlying "machinery". Our visual achievements are underpinned by an extensive and intricate system of dedicated brain areas. Shaped by our evolution, this system suggests that humans, like other primates, are "made for" vision. Brain areas dedicated to vision can be found both in the most recently evolved structures of the brain, the neocortex (cortical structures) and older structures that are enveloped by neocortex (subcortical structures). So-called "early" or "lower" visual cortical areas receive information sent by the retina via subcortical structures.

Area V1, the first and largest among these areas, will be the focus of this thesis. Early areas are thought to "compute" relatively simple features and optimize inputs to mid-level and late visual areas, for example by removing statistical redundancies (e.g. homogeneous parts of an image). Mid-level visual areas compute more complex features, and provide inputs to the final, "higher" stages of the visual system that serve more specific goals such as face recognition or visually guided grasping. The areas in this "hierarchy" are highly interconnected: connections from higher back to lower areas, and connections within an area, far exceed those in the direction from low to high (see section 1.3, Felleman and Van Essen 1991; Markov et al. 2014; DiCarlo et al. 2012; Serre et al. 2005; Salin and Bullier 1995). As a consequence of this anatomical structure, the visual system (and the brain in general) needs to solve an impressive coordination problem in order to function.

The basic "computational unit" of the brain is a highly specialized cell type, the neuron. Neurons are cells that "communicate" with each other. They come in a variety of shapes, but generally receive inputs from many other neurons, and also send outputs to many other neurons. Consequently, they look a little bit like trees, with branches on the top (dendrites) that receive information, a cell body (soma) that can also receive information, and branches on the bottom (roots in the tree analogy) that send information (axons). For a receiving neuron, the electro-chemical impulses arriving both at the soma and dendrites affect its membrane potential (the difference in electric potential between the inside and outside of the cell, a basic cell property in animals). Given appropriate conditions, sufficient depolarization of the neuron's membrane potential generates an output in the form of an electrical impulse (action potential or "spike"). The action potential travels along the axon, and in turn generates electro-chemical impulses in targeted neurons[1].

Neurons that can communicate with each other through their connectivity form a network. The computational power of the network results from the concerted behavior of its cells, and its capacity to exhibit many different concerted behaviors (Hinton, 2000; Maass, 2016). Since each neuron is influenced by many others, there are many ways to generate an action potential in a given cell. Furthermore, many cells can be active at the same time. This means that an individual spike is only meaningful in the context of the rest of the network activity: information processing is a collective effort. Crucially, interactions between neurons need to be well timed. Because individual cells integrate their inputs only within relatively short time windows on the order of milliseconds (and the world can change within milliseconds), relevant information (related inputs) needs to arrive coordinated in time (Azouz and Gray, 2000, 2003; Abeles, 1982; Softky, 1994; Salinas and Sejnowski, 2001; König et al., 1996; Koch et al., 1996; Singer, 1999; Fries, 2005). (This is about the degree to which researchers that stress integration and rate codes, and researchers that stress coincidence detection and timing, will agree. See General Discussion section 6.5.) For example, different V1 cells may receive related information that requires integration from different parts of the same objects. They may also receive redundant information that should be removed for an energy-efficient response.

How do cells, both within and between areas, coordinate their activity in time? The optimization of this coordination is a central (or possibly the central) learning goal of the brain. As a consequence, this coordination may be optimized on several

[1] The systematic study of neurons started in the late 19th century, with seminal work by Golgi and Cajal (Cajal, 1894; Golgi, 1875). For a historic overview, see DeFelipe (2015), for a general introduction, see Levitan and Kaczmarek (2015).

timescales and using several strategies. During development, a large number of connections between cells that are present at birth will be removed in a use-dependent manner. Connections that show the best coordination with other cells survive the so-called "pruning" stage (Rakic, 1994; Scelfo and Buffelli, 2009). Another related and life-long strategy the brain employs is the strengthening of connections that have useful relationships in time (e.g. the famous "fire-together-wire-together" principle of Hebbian learning, Caporale and Dan 2008; Hebb 1949). It also appears that cells adjust the conduction delays of their outputs: the branching axonal outputs of the same cell arrive in their various destinations at the same time, despite strong differences in the overall distance of the targets (isochronicity, Salami et al. 2003; Kimura 2009).

While powerful, all these strategies rest on structural changes. Structural changes take time to implement and are no match for the timescales of behavioral flexibility. For example, we can switch our focus of attention between simultaneous conversations at a party. Operating on faster timescales, another solution to the coordination problem may be to make use of the natural propensity of neural circuits to oscillate. Neural circuits in a wide range of species from lobsters to pigeons to primates show oscillatory activity (Buzsáki et al., 2013; Buzsáki and Draguhn, 2004). This rhythmic activity occurs in a broad range of frequencies (speeds) and in various functional systems, including the brainstem (breathing), motor system (walking, swimming or flying) and also in the visual system (Buzsáki, 2006). Oscillatory synchrony may enable the moment-by-moment coordination necessary for behavioral flexibility (Buzsáki and Draguhn, 2004; Fries, 2005).

In primate V1, oscillations in the gamma-frequency range (ca. 30-90 Hz, see section Gamma-band synchronization in V1) are a prominent feature. Interestingly, these oscillations are highly dependent on both visual stimulation and task or state (Buzsáki and Wang, 2012; Singer and Gray, 1995; Engel et al., 2001). Coordination in this frequency range is also interesting because it matches the time window within which different inputs need to arrive to sum effectively and drive a neuron to spike (Lisman, 2005; Börgers and Kopell, 2003; Harris et al., 2003). Consequently, it is widely believed that neurons that coordinate their activity in this frequency range will have a strong impact on a target cell. Therein lies power: such coordination may be crucial for reliable communication, for selecting competing inputs as in the case of selective attention, and may also enable Hebbian or generally timing-based plasticity (Buzsáki and Wang, 2012; Singer, 1999; Fries, 2009, 2005; Traub et al., 1998; Axmacher et al., 2006; Singer, 2008). A better understanding of V1 gamma oscillations may therefore be key to understand how V1, and by extension possibly the brain, solves the coordination problem, and is a main focus of this thesis.

Whereas the stimulus- and task-dependence of gamma oscillations is relatively well known, changes in gamma oscillations both resulting from or resulting in plasticity are much less studied. This is, however, a highly relevant question: The coordination problem must be solved moment-by-moment, but is further amplified by the ongoing need for the brain to learn and adapt. For example, we will meet new people that we need to remember, and acquire expertise with objects we frequently encounter. Whereas these are long-lasting changes, recent experiences with objects can also have more transient effects. For example, we experience familiarity for things we recently saw, discriminate better between similar objects, and we recognize an object faster if recently seen (Kohn, 2007; Brady et al., 2008; Standing, 1973; Árni and Campana, 2010). On the level of V1, which provides inputs to the rest of the visual hierarchy, responses to visual stimuli encountered repeatedly may be continually refined both transiently and long-term (Lazar et al., 2018; Kohn, 2007; De Weerd et al., 2006;

Schoups et al., 2001; Fournier et al., 2011).

Does gamma activity play a role in this? In Chapter 2, this question will be pursued by considering how visual responses in V1 change with stimulus repetition. Repetition effects are a simple form of plasticity that may occur in early visual areas, within seconds and during natural vision without instruction or task. This question will be explored with a focus on colored natural images as stimuli, as an attempt to study stimulus specificity of repetition effects using responses to naturalistic inputs. In Chapter 3, the stimulus location specificity and persistence of repetition effects on gamma-band activity will be explored.

Following the discovery of the stimulus specificity of both, repetition effects on gamma-band activity and the average gamma-band activity in V1 itself, I will then consider underlying mechanisms for stimulus specificity of gamma-band responses. I will consider both spatial structure and chromatic aspects of the stimuli. Understanding the stimulus specificity of gamma-band activity is both an important goal in its own right, and an important step to better understand mechanisms of stimulus-specific repetition effects.

To summarise, the following following questions will be considered:

In Chapters *2* and *3*,

1. What are the effects of stimulus repetition on V1 neuronal responses, especially gamma-band activity?
2. Are repetition effects stimulus specific?
3. Are repetition effects specific to the stimulus location?
4. Do repetition effects show persistence (memory)?

In Chapters *4* and *5*,

1. What is the role of color for generating gamma oscillations?
2. What is the role of spatial predictability of an image for generating gamma oscillations?
3. Do gamma-band responses to colored stimuli depend on context?

Compared to other mammals, humans share a long evolutionary history with other primates, as well as similarities in lifestyle such as social structure and diet (Osorio and Vorobyev, 1996; Changizi et al., 2006; Shepherd et al., 2010; Bernstein, 2005; Kumar and Hedges, 1998). Consequently, our visual system is very similar to that of other primates, which provide excellent model organisms to understand vision (Felleman and Van Essen, 1991; van Essen, 2004; Orban et al., 2004; Felleman, 2004; Gagin et al., 2014; Vázquez et al., 2017; Vogels and Orban, 1990; De Valois et al., 1974b,a; Mantini et al., 2012; Tootell et al., 2003). For this thesis, visual responses in rhesus macaques were studied. In the following sections, relevant concepts regarding the primate visual system will be introduced that recur throughout the chapters. Namely, the general structure and function of the visual system, and the early brain area that was the focus of study, V1, will be briefly presented. For V1, color processing and spatial contextual modulation will be addressed, followed by an introduction to gamma-band oscillations in V1. More specific concepts, hypotheses and links between gamma oscillations and stimulus repetition, color and spatial context, will be addressed in chapter-specific introductions.

1.3 The primate visual system

"Nothing is easier than to familiarize one's self with the mammalian brain. Get a sheep's head, a small saw, a chisel ..."

William James, The Principles of Psychology (1890), footnote to chapter 1.

The brain is a peculiar organ, on the one hand a part of an organism with biological properties much like other organs, on the other hand handling the very special task of generating intelligent, adaptive behavior: it is both mind and matter. The primate visual system, like the rest of the brain, can be studied at different "levels of analysis" as famously described by David Marr (1982):

1. Computational level: what problems does the system solve, and why? (e.g. face recognition to interact with conspecifics, color discrimination to find ripe fruit)
2. Algorithmic level: what processes are used, what is the language (code) to achieve the computational goals? (e.g there may be basic operations like filtering, decorrelation, or averaging)
3. Implementational/physical level: how are algorithms implemented in brain circuitry (e.g. how do neurons need to be connected to compute an average?)

These levels are interfacing in such a way that describing processes in the brain concisely yet accurately is notoriously challenging. For example, on the physical level, the retina sends nerve pulses along the optic nerve, which on the other hand is a specific pattern or "information" on the computational level. For the primate visual systen, we will begin with the large-scale implementational level, i.e. anatomy, because it is in some ways the easiest to measure, with challenges residing in the technological rather than philosophical realm.

Vision in its simplest form only requires photoreceptors (i.e. elements of a cell or cells that respond to light) whose responses are used to alter the behavior of an organism. This is already the case for single-celled algae that perform luminance change detection (Foster et al., 1984). In animals with central nervous systems, the activation of photoreceptors is still the first step, but followed by activity in specialized cells and circuits of various form and complexity. A mammalian brain exhibits structure (organization) on 6 orders of magnitude, nested within each other (Lichtman and Denk, 2011). Structure spans from subcellular structures like dendrites, to different cell types and morphologies, forming clusters and distinct brain areas. Different brain areas are distinguished by their cytoarchitecture, functional responses, and in case of vision, typically contain a representation of the entire visual field (Gattass et al., 2015, 1981, 1987; Felleman, 2004; Brodmann, 1909). On the largest scale of anatomical structure, these brain areas form a densely interconnected network, sometimes referred to as the "visual hierarchy" (Felleman and Van Essen, 1991; Markov et al., 2014; van Essen and Maunsell, 1983; Rockland and Pandya, 1979). On this level, primates show great similarity between species (in particular between larger old-world monkeys and humans), and possibly the greatest divergence to other mammals (Laramée and Boire, 2015; DiCarlo et al., 2012; Felleman and Van Essen, 1991; Tootell et al., 2003; Mantini et al., 2012).

1.3.1 Visual hierarchy, yet distributed processing

"There is no a priori reason to restrict the notion of hierarchical processing to a strictly serial sequence. In general, any scheme in which there are well-defined levels of processing can be considered hierarchical." Felleman and Van Essen (1991)

Primate vision begins with the stimulation of photoreceptors of the retina in the eyes. The retina possesses a complex machinery of its own and is involved, for example, in compression of the initial information and edge detection (Dowling, 1970; Masland, 2012; Meister and Berry, 1999). From there, two pathways are sending visual information: information processing is parallel. The thalamo-cortical route, the pathway that involves the vast majority of fibers in primates, is also the most direct pathway to the visual cortex, starting in V1 (van Essen et al., 1992; Kaas and Huerta, 1988). Information is sent and transformed along the lateral geniculate nucleus of the thalamus on to V1 and from there to higher brain areas (see Figure 1.1).

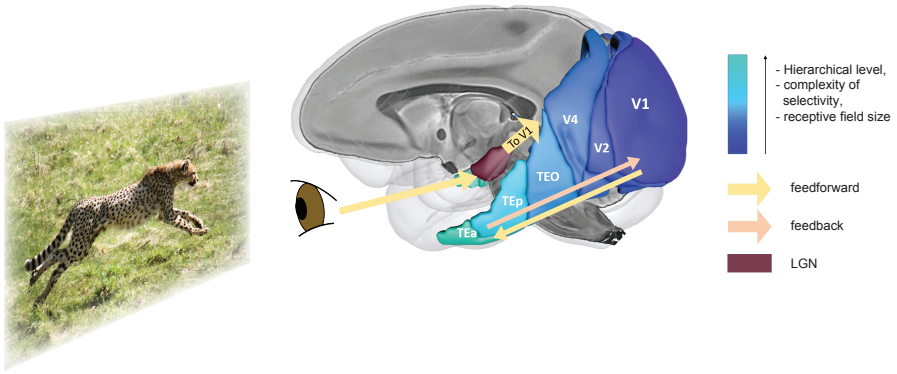

Figure 1.1 — Illustration of processing along the ventral ("what") visual stream. Images from the outside world enter the eyes as patterned light (left). The retina in the eyes processes the image, and sends output to LGN (red) and SC (not shown, see text). From there, inputs are sent all the way to the back of the brain to V1, and on along the visual hierarchy (yellow arrows, higher areas in lighter blues, see legend). Feedback information travels in the opposite direction. Note that for simplicity, not all ventral stream areas are shown, neither is the dorsal stream. The brain is seen from the side ("sagittal" view) and is that of a rhesus macaque. Other primates, including humans, have a similar organisation. Eye is not to scale. Abbreviations: V1, V2, V4 (visual area 1/2/4), LGN (lateral geniculate nucleus), SC (superior colliculi), TEO (temporal-occipital area), TEp and TEa (posterior and anterior temporal area). Image credit: photograph of cheetah adopted from Malene Thyssen, distributed under a CC BY-SA 3.0 license. 3D rendering of brain with parcellation of ventral stream areas custom-made based on Calabrese et al. (2015) for the MRI image, Paxinos et al. (2009) brain regions and the Scalable Brain atlas composer tool (Rolf, 2011).

Visual information is also processed along a second route via the superior colliculi (SC). This route is dominant in birds and reptiles, and plays a major role in vision in tree shrews (an order closely related to primates, Petry and Bickford 2019). In primates, it receives less inputs and is considered secondary (Kaas and Huerta, 1988; van Essen et al., 1992). Nevertheless, it is involved in important functions: The SC guide eye movements and gaze direction, also using direct retinal inputs. Whereas the SC do not appear to be involved in the computation of object recognition or color processing, they receive such and other information from visual and nonvisual cortex,

likely instructing movements, as the direct visual inputs do. The SC also project to the visual cortex via the thalamus (Kaas and Huerta, 1988; van Essen et al., 1992). Why have two pathways that then end up instructing movements in the SC? The answer may be that this enables fast responses based on basic stimulus properties, as well as slower, but more informed responses based on the full cortical visual machinery (Yoshida and Isa, 2011; Sparks et al., 2000; Edelman and Keller, 1996).

In the thalamo-cortical route, visual activity then cascades from V1 through the so-called "visual hierarchy" (see Figure 1.1). This happens along two major streams, the so-called "dorsal" or "how" stream of vision, involved in computations guiding motor interactions with the visual world, and along the "ventral" or "what" stream, involved in visual recognition (Goodale and Milner, 1992; Mishkin et al., 1983). While this division into dorsal and ventral provides some helpful intuitions, it should be noted that these pathways involve neighboring and often highly interconnected areas. Inputs are sent both in the direction from lower to higher areas ("feedforward"), and in the other direction ("feedback"). Furthermore, neurons have strong connectivity with areas on the same level of hierarchy and with other cells within their own respective area. Consequently, visual activity spreads rapidly along the visual hierarchy, such that different areas are engaged simultaneously, resulting in distributed processing (Felleman and Van Essen, 1991; Markov et al., 2014; Lamme and Roelfsema, 2000; Serre et al., 2005; Salin and Bullier, 1995).

Areas along the hierarchy show a gradual increase in the complexity of the visual stimuli that their cells are maximally responsive to, as well as an increase in receptive field size (Güçlü and van Gerven, 2015; Smith, 2002; Zeki, 1978; DiCarlo et al., 2012; Hubel and Wiesel, 1962). A receptive field (RF), is the region of the visual field where a stimulus elicits a response (see Figure 1.2). RF size increases along the hierarchy because of increasing convergence of inputs from different parts of the visual field. The hierarchical relationship between areas is also reflected in the termination of feedforward inputs more strongly in superficial, and feedback inputs more strongly in deep parts (layers) of each cortical area, an effect that becomes more pronounced with hierarchical distance between areas (Markov et al., 2014; Salin and Bullier, 1995).

Before turning to a more detailed description of V1, the gateway to both ventral and dorsal stream, we will consider to what degree we can expect specialized function versus general coding principles in different (visual) cortical areas. How special can we expect V1 to be? In other words, we require some more perspective on the algorithmic and computational level.

1.3.2 Current theory on sensory processing: specialization or general principles?

"Computational theory is an intermediate level of abstraction between the underlying mechanisms [...], and physiology and behavior [...]. The field of neuroscience might benefit from the recognition, in other fields of science, that reductionism is inherently limited, and that there are fundamental organizing principles at intermediate levels." Heeger (2017)

Currently, there is no consensus in neuroscience about a unifying theory of cortical function, though a few proposals have been put forward (Marr and Brindley, 1970; Grossberg, 2013, 1997; Grossberg and Pearson, 2008; Friston, 2005, 2012; Heeger, 2017; Valiant, 2014). The functional specialization of brain areas and highly specific responses of individual neurons have inspired the ever-more detailed and sophisti-

cated study of circuits serving a dedicated function, with attempts to build a theory for a specific brain area (e.g. mid-level visual area V4, Roe et al. 2012) or a particular functional goal (e.g. a color system, Livingstone and Hubel 1984; Chang et al. 2017). Much work is devoted to describe "orientation cells", "head direction cells", all the way to "grandmother cells" (famously referring to a cell responding whenever a person sees their grandmother, Gross 2002; Thomas and French 2017). This approach emphasizes distinct computations in different circuits. This perspective is certainly not wrong in the sense that vision (and the brain) has several distinct goals, like recognizing a face or catching a ball.

On the other hand, it is fairly clear that there are more or less similar neurons with more or less similar circuit motifs repeated over and over in visual cortex: whatever the function of an area, the building blocks are the same (so-called "canonical microcircuits", Douglas and Martin 1991; Douglas et al. 2008; Bastos et al. 2012; Fries et al. 2012, but see van Hooser 2007). Differences must originate in the inputs (e.g. sensory modality) and goals that may direct plasticity. Further, it has been suggested that these physical circuits implement canonical computations on the algorithmic level (e.g. Carandini and Heeger 2012; Dasgupta et al. 2017). Such a hierarchical repetition of a few canonical operations is at the core of current machine vision (DiCarlo et al., 2012).

Given such standardized building blocks, what are overarching general principles and goals that need to be achieved? Adaptive behavior requires an appropriate model of the world, in more sophisticated cases including oneself and other agents. A working model can be used to make inferences: given a combination of sensory input and current context and expectations, what is currently most likely going on in the outside world? Such a model can explore different possible states of the world (finding different alternative explanations for current sensory input, recall memories, generate ideas) to make those inferences. It also makes predictions (inference about the future, what is the future state of the world?). Note that given the delay between events in the outside world and neuronal responses, prediction is a constant, rather than occasional, task for the visual system (Nijhawan, 2008). Simultaneously to inference and prediction, any implementation of these principle goals will also need to allow for continuous learning. David Heeger (2017) recently put forward modeling work suggesting that inference, prediction, exploration and learning can be achieved in the same network depending on its current "state". Given state parameters, the network weighs sensory inputs, prior expectations, and context differently in different states. Previous theoretical and empirical work has also put a strong emphasis on the integration of sensory evidence and expectation (e.g. Grossberg 2013; Friston 2005; Ma et al. 2006; Körding and Wolpert 2004) and increasingly prediction of the future (e.g. Palmer et al. 2015; Singer 2018; Wacongne et al. 2012; Hawkins et al. 2009)[2]. There is therefore some agreement on these computational goals. Interestingly, different theories make different predictions about the algorithmic and neuronal implementation of these goals, with consequences including predictions about the mesoscopic neural level like gamma oscillations. These will be considered in more detail in Chapter 4.

First however, we should note that there is a second set of goals to fulfill. This is because the brain is biological organ and not an idealized computational machine. Neural activity comes at a cost. Since there is a cost to both neural activity but also

[2]Prediction of the future may sound redundant. However, as pointed out by Heeger (2017), a lot of theoretical work uses the word "prediction" or "predictive coding" describing a process of predicting the current input from expectations and the past, rather than predicting the future.

to errors in behavior, there will be complementary goals, some of which may require a tradeoff. Specifically, an optimal code will be both efficient and reliable. Efficiency here means that as little spiking (or generally neural) activity should be used as possible, to reduce metabolic cost. Reliability means that the goal of a specific computation will be achieved repeatedly in the face of noise from the outside world and from neural activity. A major way to improve reliability is through redundancy. Reliability and efficiency may be at odds with one another - although this is a central question in neuroscience where neither theory nor evidence are fully developed (Pryluk et al., 2019; Barlow, 1959, 2001; Kohn et al., 2016; Chang et al., 2017)[3]. Another goal of the brain as a physiological organ must be homeostasis: many elements in the brain need to remain at a physiological equilibrium (Tononi, 2009). This will for example limit neuronal excitability levels to non-epileptic states, regardless of theoretically optimal ranges of firing responses. It is currently largely unknown how the brain solves tradeoffs between different goals, under which circumstances the tradeoffs exist and if the brain may use different strategies in different brain areas.

When studying V1, it is both interesting to ask what the unique contributions or properties of this area may be, but also if we can infer any relationships to general principles and goals. Given that a theory attempts to cover all of cortical function, finding counterexamples even in just one area can be informative. We will explore concepts of inference, prediction, efficiency and homeostasis further when considering why V1, and particularly V1 gamma-band activity, may behave the way it does (section 1.4.2, Chapters 2, 3, and 4). First we will consider the overall layout of V1.

1.4 Area V1

"Given how little of V1 function we can currently claim to understand, we should be prepared for some surprises as new data come in." In: "Do we know what the early visual system does?" Carandini et al. (2005)

V1 is a large cortical area at the bottom of the visual hierarchy. It is also called primary visual cortex, or due to its special cytoarchitecture, "striate cortex". Animals (and in particular primates) with high visual acuity have a large V1 with small receptive fields (Nienborg, 2004; Baker, 2013; Veilleux and Kirk, 2014; Srinivasan et al., 2015). Tiny primates with high visual acuity devote an unusually large percentage of their brain to V1, in line with the idea that visual acuity is at odds with compression (Ho et al., 2019). Receptive fields in V1 are substantially smaller than in later parts of the hierarchy (Harvey and Dumoulin, 2011; Smith, 2002; Zeki, 1978). Neurons with similar spatial RF locations are nearby in cortex, such that the visual field is tiled into a retinotopic map in V1 (see Figure 1.2).

This mapping is not 1:1 with the real world, however. The foveal region of the visual field (the central 2-5 degrees of visual angle, dva, in the visual field) is massively overrepresented. The foveal region is the central point fixated by the eyes, where visual acuity is highest due to the least compression in the retina: Whereas peripheral photoreceptors are pooled over a larger region of the visual field, foveal photoreceptors in the retina are very dense, and are mapped 1:1 (minimally) onto the output

[3]Reliability may also be increased through checks and balances via reafferent information or feedback. In general, degeneracy can increase reliability. Degeneracy in a system means that different elements can affect the output in a similar way, but crucially, at the same time can have independent effects (Tononi et al., 2002). An example we encountered earlier is the thalamocortical versus retino-tectal pathway, both affecting the SC.

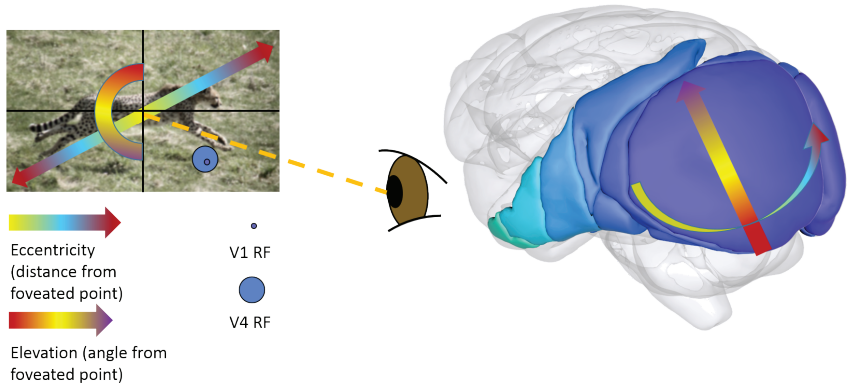

Figure 1.2 — Illustration of retinotopy in V1. Left: illustration of eccentricity and angle of a point of an image with respect to the foveated point (endpoint of orange dotted line). In blue, the receptive field (RF) coverage of example neurons from V1 (early) and V4 (higher up the hierarchy) are shown. Right: Brain of Figure 1.1 seen more from the back, to illustrate the size of V1 and sketch its retinotopic organization. The neurons with foveal RFs (or low eccentricity) are located at the most lateral part of V1 (yellow part of bent arrow), the most eccentric RFs are going into the sulcus (cortical inward folding) and below the part of V1 visible from the surface. Note that a disproportionate amount of V1 is devoted to the fovea, or central vision. Neurons are also arranged such neurons with similar visual angles are close to each other (straight arrow). Note the inversion of the representation: the lower visual field (below fixation) is on the dorsal (upper) part of V1. Image credit: see Figure 1.1.

neurons of the retina (Wässle and Boycott, 1991). As a consequence, RF size also changes with eccentricity, with the smallest RF at the fovea being well below 1 dva in size.[4] Due to the lack of compression, foveal regions encompass an extensive portion of V1 (Wässle et al., 1989). This retinotopic map with a magnified fovea, where nearby points change smoothly in eccentricity and elevation, has a third dimension: the cortical layers.

Like other neocortical areas, V1 has a laminar structure with 6 layers from the outermost layer (1) to the innermost layer (6) adjacent to the white matter (i.e. axonal fiber tracts, which are fatty and as a result appear white). Neurons across layers respond to the same part of the visual field (see Figure 1.3).

Inputs from the thalamus arrive predominantly in so-called layer 4 (or due to its cytoarchitecture, granular layer, roughly in the middle), as well as layers 5/6. Using laminar electrophysiological recordings, this can be detected with current source density (CSD) analysis (e.g. Mitzdorf 1985; Maier et al. 2010, Lowet, Roberts, *Peter*, Gips and De Weerd (2017), see Figure 1.4). Figure 1.4 illustrates that whereas RF location is constant across cortical layers, other features are not. For example, a prominent source of gamma oscillations resides in superficial layers 2/3.

Layer 4 strongly projects to layers 2/3, which shows more complex and invariant response properties than layer 4 (see below for an example of invariance). Layers 2/3 are the primary output layers to the downstream visual hierarchy. Information also flows from layers 2/3 and 4 to the deeper layers, which receive top-down feedback and communicate with the motor system. That is, even in V1, as at every other hierarchical level, there is an output that affects behavior directly and cuts out of the feedforward

[4] One degree of visual angle, or dva, is about the size of the fingernail of your thumb if you hold your arm outstretched right in front of you.

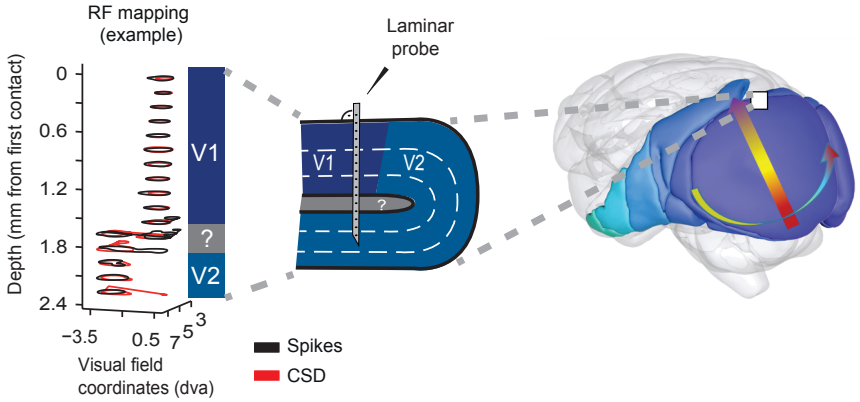

Figure 1.3 — Example receptive field mapping with a laminar array of electrodes. Right: view of the brain as in Figure 1.2. White square indicates approximate recording location. Middle: sketch of cortical folding of V1 into V2 with main laminar compartments, including sketch of electrode position. Electrode size is approx. to scale. Left: View of receptive fields along the array. Note a distinct jump in RF position with the switch of the visual area. RFs were measured using a sparse-noise procedure using both spiking responses and current source density (CSD), the second spatial derivative of the local field potential (see Lowet et al. 2017 for methods details).

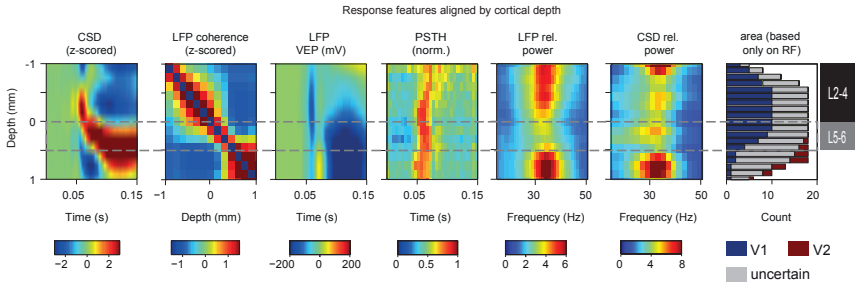

Figure 1.4 — Laminar view of some basic electrophysiological features in response to visual stimulation. Data represents cross-session aligned average from an example animal (M2 of Lowet et al. 2017). Grey dashed lines indicate location of layers 5-6 based on the CSD reversal point (leftmost panel). Next to CSD, the LFP coherence between all contact pairs is shown, which has been proposed as an alternative for finding the layer 4/5 border (Maier et al., 2010). Next to this is the visual evoked potential (VEP), which can also be recorded with single electrodes. Note the lack of spatial resolution in the response pattern compared to CSD. This means that the laminar position of electrodes is difficult to assert based on single electrodes in awake animals. The peri-stimulus time histogram (PSTH) of each session was normalized to the maximum activity of the maximally active channel. Relative power was computed as log(stimulus/baseline) for both local field potential (LFP) and CSD. Rightmost column shows the number of contacts assigned to each depth and their assignment to V1 vs white matter or V2 based on receptive field mapping alone, providing an estimate independent of the CSD reversal point.

cascade. Layer 1 contains mostly dendritic processes and some inhibitory neurons, making this layer exceedingly difficult to study using awake primate physiology. Recent rodent work suggests a role for these neurons in contextual gating of learning (Letzkus et al., 2011).

What does V1 do? In primates, V1 cells are the first cells capable of detecting oriented edges (in other species this can happen earlier). This is a crucial step in vision: edges and their orientation provide an important part of the object information contained in an image (Marr, 1982). Some cells display invariance to the direction of the brightness difference that defines the oriented edge (the defining feature of complex compared to simple cells). The development of complex cells is an example of a feature invariance computation that can be readily understood by a selective summation of inputs (Hubel and Wiesel, 1962; Priebe and Ferster, 2012; Movshon et al., 1978)[5]. First, inputs from LGN cells with receptive fields that are aligned along a particular direction in the visual field are pooled, yielding a "simple cell". A simple cell has subregions in the RF that will only respond to a bright stimulus, which are neighboring subregions that respond to a dark stimulus, defining the edge. Correct integration of several simple cell responses (corresponding to an OR computation) yields complex cells. Note that this is an entirely feedforward computation where inputs are integrated from one stage to the next. Given the prevalence of simple cells in layer 4 and complex cells in other V1 layers, it seems clear that this computation is achieved in V1 (Yu and Ferster, 2013; Martinez et al., 2005).

Next to orientation, V1 cells seem to be detectors for other features like higher-order spatial correlations and binocular depth (Yu et al., 2015; Cumming and DeAngelis, 2001). They typically respond within a certain range of contrast, motion direction and spatial and temporal frequency, which form orthogonal axes with which to describe visual images ("frequency space", Mante and Carandini 2005). Collectively, neurons in V1 are thought to tile this frequency space, just like they tile the visual field. In addition, V1 is highly color selective (see section 1.4.1). Given the tiling of the visual field and the "tuning" (combined response preference for various visual features) of V1 cells, one can think of V1 as an array of feature detectors that are repeated across the retinotopic map. This conceptualization is a major inspiration for current artificial neural networks (ANNs, DiCarlo et al. 2012). However, this view is incomplete: For example, V1 responses are modulated by task requirements and show evidence for inference of missing information, in particular for edges (Kapadia et al., 1995, 1999; Grosof et al., 1993; von der Heydt et al., 1984; Roberts et al., 2007; Chalk et al., 2010). This inference relies on integration of information beyond the RF of any given V1 cell. Such integration is a crucial feature of V1, likely related to gamma activity and will be introduced in section 1.4.2.

We will first briefly consider color processing, as it is an important aspect of natural vision and plays a role in Chapters 2, 4, and especially 5. We then turn to center-surround integration mechanisms, a core feature of neocortical circuits. Finally, I give a brief overview of gamma oscillatory activity in V1.

[5] Whether this is how it is actually achieved, and whether invariance is the actual goal, remains to be established.

1.4.1 Color in V1

"Colour is what the eyes see best."

Title of paper by Chaparro et al., Nature, (1993), comparing luminance to color discrimination.

A distinguishing feature of old-world primates compared to other mammals and most new-world primate species is trichromatic vision (Surridge et al., 2003). Trichromacy in primates correlates with frugivory (i.e. a fruit focused diet, Osorio and Vorobyev 1996). Notably, trichromatic individuals enjoy advantages in food retrieval in species that have both dichromatic and trichromatic members (Melin et al., 2017). In some primates, including the rhesus macaques studied here, color also plays an important role as a social cue (Waitt et al., 2006; Gerald et al., 2007). Trichromacy is based on three specialized receptors in the retina (so-called cones) with sensitivities for different parts of the wavelength spectrum of visible light. A fourth type of receptor, the rods, is most responsive under dim light conditions. In contrast, cones are responsive under daylight brightness conditions (Wässle and Boycott, 1991). L-cones respond to the *l*ongest wavelengths, M-cones to inter*m*ediate wavelengths, and S-cones to *s*hort wavelengths. Color vision is based on response differences between these cones (color opponencies). This is likely because the absolute value in responsiveness of any single cone type depends on the overall level of brightness and is therefore uninformative about color on its own. Response differences are computed already on the level of the retina and LGN and include L-M ("red-green") contrasts as well as S-LM ("yellow-blue") contrasts. Red-green contrasts cannot be computed by dichromats and are therefore essential for the trichromatic behavioral advantages. Yellow-blue contrast responsive cells appear to be the minority of cells in both retina and LGN, where L-M responses dominate (Wässle and Boycott, 1991; Shapley and Hawken, 2011; De Valois, 1965). Nevertheless, the blue-yellow contribution may be recovered in V1, where it appears considerably stronger than in its inputs (De Valois et al., 2002). A do-it-yourself demonstration of color opponencies can be enjoyed by prolonged fixation of a colored surface, followed by fixation on a white surface (for example the dark blue surface of V1 in Figure 1.2, although red or green surfaces may yield stronger results). Due to fatigue in one color channel, the other cone response will dominate and the result is a visual afterimage in the color opponent to that fixated.

Early studies have noted a distinct cytoarchitectonic patterning in V1 when staining with cytochrome oxidase, yielding cytochrome oxidase rich "blobs" and between them, "interblobs". It was found that blobs contain very high numbers of color-selective cells, whereas interblobs were suggested to be more selective to orientation (Livingstone and Hubel, 1984). Blobs and interblobs project onto specific parts of the downstream area V2. This suggested that there are distinct color and shape systems in V1 and subsequently in the visual hierarchy (Shapley and Hawken, 2011). Indeed, there are some specific subregions of higher-order visual cortex that appear to be very strongly responsive to colored stimuli (Chang et al., 2017; Conway and Tsao, 2006). Despite this apparent specialization, more recent studies using a wider array of visual stimulus types have reported that V1 is highly color selective throughout both blob and interblob regions (Shapley and Hawken, 2011; Schluppeck and Engel, 2002; Wachtler et al., 2003).

What kind of computational function may V1 serve with respect to color? There is currently no consensus on an answer to this question. One possibility is that V1

provides a rotation of the "axes" of the color system from L-M and S-LM to axes that match our perception (Lafer-Sousa et al., 2012; Wachtler et al., 2003; Horwitz and Hass, 2012). Given the larger diversity of responses in V1 compared to its inputs from LGN (De Valois et al., 2002; Wachtler et al., 2003), it may also be used as a larger set of responses that higher order areas can draw on to support color computations. Wachtler et al. 2003 make a link to general computational principles, suggesting that the transformation from LGN to V1 corresponds to an implementation of an efficient coding of natural image statistics.

One robust finding is that a particular new cell type emerges, called a double-opponent cell (Shapley and Hawken, 2011). These cells are responsive to color-defined edges rather than surfaces, whereas LGN cells respond most strongly to surfaces or spots. In analogy to complex cells for orientation, these cells are also emerging in the upper layers of V1. The predominance of color-selective cells in V1 attests to the importance of color information for trichromatic primates like humans. Another line of evidence comes from the simple comparison of achromatic to chromatic (colored) stimuli. Compared to achromatic (black-white) flashes of light, chromatic flashes induce an additional response in V1 CSD responses. A similar chromatic advantage is seen across the ventral visual stream (Chen et al., 2007). Given the foveal overrepresentation in V1 and the cone density of the fovea, this is not surprising but emphasizes the strong role of color in natural vision for primates. Color responses in V1 have been studied mostly on the level of single neuron tuning, and using fMRI. In constrast, and inspite of the known stimulus-dependency of gamma oscillations for other feature dimensions, mesoscopic signal like oscillatory activity have received very little attention. Color-specific responses of gamma oscillations will be the topic of Chapters 4 and 5.

1.4.2 Contextual modulation in V1

"The challenge facing the visual system is to extract the "meaning" of the image by decomposing it into its environmental causes. For each local region of the image, that extraction of meaning is only possible if information from other regions is taken into account." Albright and Stoner (2002)

Whereas trichromatic color processing concerns itself with a specifically visual feature, the integration of feedforward inputs with contextual information may be a fundamental mechanism that spans brain areas, species and sensory modalities (Spillmann et al., 2015; Nurminen and Angelucci, 2014; Sutter et al., 1999; Olsen and Wilson, 2008; Maffei and Fiorentini, 1976; Gilbert, 1992, 2013; Coen-Cagli et al., 2015; Vega-Bermudez and Johnson, 1999; Hubel and Wiesel, 1965; van Den Bergh et al., 2010; Sun et al., 2002; Albright and Stoner, 2002). In the visual system, it is present already in the retina and the main input structure to V1, the LGN (Solomon, 2006; Alitto and Usrey, 2008). From the perspective of a V1 neuron with a small receptive field, spatial context is provided by neighboring neurons with neighboring receptive field locations, and feedback from neurons with larger receptive fields: it is modulated by its surround. Surround modulation has been suggested to constitute a form of gain control. Gain control is a central mechanism in sensory systems to deal with the high range of input variation. For example, luminance levels vary by around 9-10 orders of magnitude between bright daylight and nighttime vision, and cones exhibit gain control (Korenbrot, 2012). In V1, surround modulation as gain control would normalize firing rates by scaling the amount of drive to the RF with the amount of drive to the

surround (Carandini and Heeger, 2012). However, gain control may not be the only mechanism at play (Coen-Cagli et al., 2015). Some surround modulation effects seem very useful. For example, surround modulation can be facilitatory for weak stimuli in the RF. Crucially, this happens only if the surround information matches. For example, a faint oriented bar in a RF will elicit a stronger response if the surround consists of a matching bar that continues the bar outside the RF (Kapadia et al., 1995). This can be interpreted as a signature of inference: the evidence from the RF for the presence of a bar is strengthened by matching evidence in the surround. In contrast to weak stimuli in the RF, responses to stronger (higher contrast) stimuli are often reduced when the RF surround is stimulated. This effect is called surround suppression (Nurminen and Angelucci, 2014; Spillmann et al., 2015). Surround suppression, like surround enhancement, is particularly pronounced if the surround information matches the RF information (Cavanaugh et al., 2002; Sillito et al., 1995; Trott and Born, 2015). Surround modulation may play an important role in different but related functions like contour integration, figure-ground segregation, or perceptual filling in (Liang et al., 2017; Wachtler et al., 2003; Zweig et al., 2015; Lamme, 1995; Li, 2002; Coen-Cagli et al., 2012). Furthermore, models of efficient and predictive coding theories have attempted to capture the various surround modulation effects based on general computational goals (Schwartz and Simoncelli, 2001; Simoncelli and Olshausen, 2001; Coen-Cagli et al., 2012, 2015; Rao and Ballard, 1999; Barlow, 1959, 2001; Vinje and Gallant, 2000; Friston, 2005; Spratling, 2010). According to efficient coding theories, surround suppression occurs to remove responses to image redundancies across space. Predictive coding theories postulate that neuronal responses reflect a comparison between inputs to the RF and predictions generated by the surround. Predictable spatial relationships will result in a lower firing response.

Surround modulation is typically studied with firing rate responses as the variable of interest in mind. However, it is also possible that surround modulation affects temporal aspects, such as correlated or synchronized responses in neuronal populations. In Chapter 4, I will review the current literature on center-surround interactions in relationship to gamma synchronization, and explore predictions made by predictive coding theories on gamma synchrony using colored stimuli.

1.4.3 Gamma-band synchronization in V1

"The specific physiological functions of brain rhythms vary from the obvious to the utterly impenetrable." Buzsáki (2006).

Rhythmic or repeating activity is common in neuronal networks: they oscillate. The ubiquity of oscillations across cortical and subcortical structures, as well as species across the animal kingdom (including insects, arthropods, reptiles, birds and mammals) suggest that the ability to generate oscillations arose early in evolution (Buzsáki et al., 2013; Buzsáki and Draguhn, 2004). Oscillations occur at a wide range of frequencies (speeds), from the very slow (taking several seconds, e.g. every 2 s or 0.5 Hz) to the very fast (repeating e.g. 100 times per second, 100 Hz). A short taxonomy of rhythms or band-limited activity commonly encountered in the brain is provided in Table 1.1.

Oscillation research has a long history, with names largely following tradition because underlying mechanisms are not fully understood. The gamma-frequency range is defined somewhat arbitrarily as 30-90 Hz following Buzsáki and Wang (2012) and Freeman (2007). Here and in the following chapters, as is typical in the literature,

Name	Freq. (Hz)	dominant state and brain regions
slow	<1	cortex & thalamus, sleep, anaesthesia
delta	1-4	cortex & thalamus, sleep, anaesthesia
theta	3-10	cortex, cerebellum & HC, wakefulness
alpha	8-13	cortex & thalamus, wakefulness, drowsiness
beta	10-30	cortex & basal ganglia, wakefulness, stim.
gamma	30-90	cortex & HC, wakefulness, stim.
epsilon/high-gamma	80-250	cortex & hippocampus, wakefulness, stim.
ripple	140-200	HC, wakefulness, stim.
fast	>250	cortex & cerebellum, sleep and wakefulness

Table 1.1 — List of some of the most common names for frequency bands of brain activity. stim. = sensory stimulation. HC = hippocampus. See Buzsáki (2006); Freeman (2007); Wang (2010); Penttonen and Buzsáki (2003).

the terms "gamma", "gamma-band activity" (GBA) and "gamma oscillations" will be used interchangeably. Gamma synchrony refers to synchronous activity between cells or areas. Gamma synchronization, strictly speaking, refers to the process with which gamma synchrony is achieved (see below), but can also be used to refer to gamma synchrony. An important distinction needs to be made between band-limited gamma-band activity and so-called high-gamma or epsilon activity (ca. 80-250 Hz). This activity occupies a broad frequency range, and has been recently dissociated from visual gamma-band activity based on pharmacological intervention and stimulus dependence (Bartoli et al., 2019; Leszczynski et al., 2019). An excellent perspective on the broader topic of brain rhythms is provided by Gyorgy Buszaki (2006).

Neural networks exhibit a range of different rhythms, and there are many ways for a network to generate any given rhythm. Two main classes of origins of oscillations are network oscillations and cell-intrinsic oscillations (Wang, 2010). The two classes will be considered in turn, with a focus on the generation of gamma-band oscillations in visual cortex. Network oscillations arise from interactions between cells, typically with opposing forces of excitation and inhibition. For example, let us consider a simplistic network of one excitatory neuron reciprocally connected to an inhibitory neuron. In response to a constant input, the excitatory neuron will fire an action potential once its membrane potential is sufficiently depolarized. This will elicit a spike in the inhibitory neuron, which in turn suppresses further spiking in the excitatory neuron. As a consequence, the inhibitory neuron loses its input and stops to suppress the excitatory neuron, which will then be able to respond again, restarting the cycle. Such a mechanism, implemented on the level of a larger population, has been suggested to underlie gamma-band synchronization (pyramidal-interneuron gamma or PING, e.g. Tiesinga and Sejnowski 2009; Börgers and Kopell 2005). Indeed, stimulation of both soma-targeting and dendrite-targeting subtypes of interneurons can generate gamma-band responses in cortex (Cardin et al., 2009; Veit et al., 2017). (A more detailed discussion of interneuron cell types, their function and gamma oscillations will follow in the General Discussion, section 6.4, see also Buszaki and Wang (2012) for review of gamma mechanisms.)

Oscillatory activity can also arise in single cells. Single cells, even in a disconnected in vitro state, can generate oscillations from non-oscillatory input through counteracting membrane currents: the oscillatory activity is inbuilt. Such cells can act as pacemakers that organize a network. They are prominent in subcortical structures

including the thalamus (Buzsáki, 2006; McCormick and Pape, 1990) but also neocortex (Sanchez-Vives and McCormick, 2000). They may also contribute to V1 gamma oscillations, where so-called "chattering cells" show gamma-range burst firing (Gray and McCormick, 1996). In addition to generating oscillations, single cells can also resonate at particular frequencies (Buzsáki, 2006). In summary, from the perspective of a single neuron, its dendritic inhibition, somatic inhibition, as well as its own physiological properties may contribute to gamma oscillations in its membrane potential or firing activity.

An additional source of difficulty for interpreting gamma-band activity is that it can be measured at various spatial scales and using various measures, which differ in sensitivity. Oscillations can occur at the level of the single cell membrane potential, single cell firing, local networks whose activity may be reflected in the local field potential (LFP), all the way up to macrosopic signals that can be measured through the intact scalp (electroencephalography, EEG, and its magnetic equivalent, MEG, Buzsáki et al. 2012). LFP and EEG signals derive from the superposition of the activity of many cells. The strength of the contribution of individual cells depends on their spatial alignment with other cells and synchrony with those cells (potentials of opposing directions can cancel each other out, Buzsáki and Wang 2012). The larger the electrodes, and the further the electrodes are away from the tissue, the more activity will be averaged across microcircuit rhythms. This means that with EEG or LFP measurements using large (mm scale) electrodes as commonly used in human epilepsy patients, the biophysics of the underlying signal may be very different from that observed using local microelectrodes several orders of magnitude smaller.

Oscillations are characterized by their frequency (speed), amplitude (strength, its square is referred to as power), and at a particular moment in time, their phase (position in the oscillatory cycle). In addition, oscillations are not necessarily sinusoidal and are further characterized by their shape. Periodic components in a signal can be detected using a Fourier analysis (named after its original developer Fourier 1822). A given signal is decomposed into sinusoids of different frequencies and their respective amplitudes and phases. This allows the description of the signal in the frequency- rather than time-domain. A power spectrum (power spectral density, PSD) shows which frequencies are represented with which strength in a signal. A white noise signal will be flat in the frequency domain: all frequencies have the same power.

EEG and LFP signals typically follow a power law, their strength decreases in a log-linear fashion with increasing frequency (Penttonen and Buzsáki, 2003). While many natural phenomena follow power laws, the reason for this power-law behavior is not understood and interpretations range from the trivial (filtered noise from volume-conducted global network activity) to the fundamental (self-organized criticality, Miller et al. 2009b; Pritchard 1992; Hesse and Gross 2014; Liu et al. 2014). The presence of peaks away from the power-law relationship is typically considered an indication of the presence of oscillations that contain a substantial part of the total power of the signal. This can be indicative of coordination among larger parts of the neuronal population.

In contrast, how is coordination between individual cells measured? Two neurons (or recording sites) can be coherent when the phases of firing with respect to an oscillation in a particular frequency band show a reliable relationship ("locking"). In visual cortex, this is typically not a constant phenomenon, but such coherence occurs intermittently. Fluctuations in the individual frequencies of each neuron can reflect mutual interactions that bring about intermittent coherence (Lowet, Roberts, *Peter* et al., (2017)). This means that it is better to think of synchronization as an ongoing

process rather than a state. Synchronization can bring about an average level of synchrony (phase coherence) over a longer period of time. Phase coherence in the spiking between two cells, or even oscillatory activity for a single cell, can be difficult to detect if a cell fires relatively rarely. This is a common phenomenon in visual cortical excitatory cells, which typically do not spike on every, or even most, cycles of a gamma oscillation ("cycle skipping", Nikolić 2009). Coherence between the spiking of a cell and the local field potential can be a very powerful tool to overcome this problem and uncover the oscillatory nature of seemingly random spikes of isolated cells (Vinck et al., 2012)[6]. Regardless of its origins and measurement scale, it is clear that the amplitude (strength) and frequency of gamma oscillations is strongly modulated by visual stimulation (a small sample of examples: Hadjipapas, Lowet, Roberts, *Peter* and De Weerd (2015), Roberts et al. 2013; Gray et al. 1989; Ray and Maunsell 2010; Brunet et al. 2015). V1 gamma oscillations that are stimulus-related strongly increase in amplitude both in layers 2/3 and 5/6, i.e. the output layers of V1 (Lowet, Roberts, *Peter*, Gips and De Weerd (2017), see Figure 1.4). Early studies found that gamma synchronization occurs when neighboring RFs are stimulated in a way that follows Gestalt principles, for example by the same visual object (Gray et al., 1989; Singer and Gray, 1995; Roelfsema et al., 2004; Singer, 1999; Milner, 1974).

The notion that neurons share a gamma rhythm when processing the same object inspired theoretical proposals about the function of gamma-band synchronization. In particular, it was thought that gamma oscillations could integrate object-specific information from disparate sources, be it different cells within an area or between areas, thereby solving the "binding problem" (Treisman, 1996) through "binding by synchrony" (BBS, Singer and Gray 1995; Singer 1999; von der Malsburg 1994). The crucial idea is that gamma activity can define cell ensembles.

A related proposal also suggests the existence of gamma-defined cell ensembles, but focuses on competition between different ensembles (for example for different objects) and interareal interactions. The central idea is that gamma-coherent interactions between brain areas may lead to selective routing of information about a particular ensemble along the visual hierarchy (Fries, 2005). Selective, prioritized processing of some stimuli versus others is considered important because it may underlie selective attention to one stimulus at a time. According to the "communication-through-coherence" (CTC) hypothesis (Fries, 2005, 2015), the prioritized ensemble will align its rhythm with the downstream, receiving area such that its inputs arrive when the neurons in the receiving areas are depolarized. Such inputs should have a strong impact on the receiving area. Some evidence for and criticism of BBS and CTC will be discussed in Chapter 6.

Grouping (BBS) and communication (CTC) hypotheses concern themselves with a situation where connections of the networks in question are stable. Another suggested function of gamma-synchronized activity is the promotion of plasticity processes. Synchrony in the millisecond range induced by gamma oscillations may facilitate spike-timing dependent plasticity (Wang, 2010). On the other hand, plasticity may induce rapid changes in gamma responses themselves. In particular, repeated presentation of a stimulus can induce changes in gamma-responses to this stimulus (Brunet et al., 2014). The potential role of gamma-band activity in learning and in particular for stimulus repetition will be explored in the next chapter.

[6]The advent of recording methods with very dense arrays, recording many nearby cells, may enable another way to detect and visualize oscillations.

Chapter 2

Stimulus specificity of repetition effects

2.1 Acknowlegdements for Chapter 2

The experiment described in this chapter reflects recordings performed in the labs of Pascal Fries, Wolf Singer and Michael Schmid, and generous support from all principal investigators. On the practical side, I am indebted to Johanna Klon-Lipok, Liane Klein, Jarrod Dowdall, Marieke Schölvinck, Katharine Shapcott, and Kleopatra Kouroupaki, for introducing me to their recording setups and animals. Preprocessing pipelines to generate MUA and LFP estimates from broadband data were developed by Joscha Schmiedt. RF mapping procedures were developed by Katharine Shapcott, Joscha Schmiedt, Cem Uran and previous members of the Singer lab (monkey A, H, K) and Jarrod Dowdall and Marieke Schölvinck (monkey T).

2.2 Introduction

The repeated encounter of a sensory stimulus typically changes how the brain responds to the stimulus. Across brain areas and paradigms, the repetition of an identical stimulus tends to lead to decreased neuronal firing rates (Solomon and Kohn, 2014; Vogels, 2016) and fMRI signal (Sawamura et al., 2005; Grill-Spector et al., 2006). Stimulus repetition effects could reflect fatigue, or alternatively, an adaptation of the visual system. For example, it could improve processing of the familiar, lead to better detection of the novel, or both (Solomon and Kohn, 2014; Wissig and Kohn, 2012; Vogels, 2016; Gotts et al., 2012). Given the frequency of repeated encounters in the sensory world, repetition effects may reflect common operations of the nervous system. More generally, effects of stimulus repetition may give insight into the computational goals of the brain by providing it with an opportunity to optimize its responses under the experimenter's watch. Repetition effects therefore have the potential to reveal general goals of visual coding, such as improved efficiency, precision, accuracy, discriminability or detection (see General Introduction, section 1.3.2).

Decreased responses with repetition could have many beneficial underlying reasons. For example, responses could become more sparse, increasing SNR through removal of irrelevant or redundant information, or become more transient. Another potential, and not mutually exclusive, mechanism would be increased synchronization of the remaining stimulus-related firing responses. This could enhance or at least maintain their impact on downstream neurons, possibly while reducing metabolic cost.

Evidence for increased synchronization was reported in a previous study from our lab. Strong changes in gamma-band synchronization within and between early visual areas were found with hundreds of repetitions of grating stimuli (Brunet et al.,

2014). Fascinatingly, putative inhibitory interneurons in particular increased gamma-band synchronization. In contrast, weakly, but not strongly stimulus-driven excitatory cells dropped out of the rhythmic activity. This is suggestive of a particular form of sparsening of the response through synchronization. However, since the stimuli in this study never changed over the course of the session, it remained unclear to what degree the observed effects were specific to a particular stimulus, or if the increased synchronization would transfer to an arbitrary different stimulus. Moreover, the grating stimuli were well-known to the animals and designed to optimally drive visual gamma-band responses. The current study therefore aimed to test for the stimulus specificity of the repetition effect, using naturalistic stimuli that were initially unfamiliar to the animals. We asked whether

1. the repetition effect is stimulus specific
2. the repetition effect will generalize to a broader range of stimuli, which are more naturalistic and less familiar.

To put the current experiment in a broader context, different theories on repetition effects, common designs of repetition experiments and some of the most pertinent experimental results will be introduced in the following sections.

2.2.1 Current theories of repetition effects

Several terms have been used to describe how the brain changes its responses with stimulus repetition. The most common terms are "adaptation" and "repetition suppression", the latter is used predominantly in the higher levels of the visual hierarchy[1]. Unfortunately, adataption refers to a wide range of phenomena with a possibly equally wide range of underlying mechanisms (Solomon and Kohn, 2014). Furthermore, adaptation can refer to phenomena where a stimulus is presented for a prolonged period of time, rather than repeatedly. Repetition suppression is an equally problematic term, as it may be used interchangeably with "adaptation" to refer to phenomena in the higher visual system, and further, it is not clear whether suppression of responses (i.e. inhibition) underlies this group of phenomena (Vogels, 2016). I will therefore use the umbrella term "repetition effect" to describe response changes with stimulus repetition and "adaptation" for prolonged stimulation unless otherwise noted (see also Discussion section 2.5.5 for a consideration of stimulus exposure duration versus repetition).

Currently, there is no commonly accepted theory that encompasses all effects of stimulus repetition, and links functional and mechanistic perspectives. Explicit computational modeling and suggested mechanisms of particular effects exist, however (see next section). Furthermore, there are some proposals about some potential general mechanisms of stimulus repetition, a selection of which will be reviewed here.

Output fatigue. One group of proposals entails the idea of fatigue, and along with it the notion that repetition effects can have detrimental effects on perception and behavior, if any. According to the "output fatigue" hypothesis, neuronal responses reduce proportionally to their amplitude. Neurons are stimulated so strongly that the mechanisms that ensure their responsiveness to inputs cannot keep up. Artificial stimulation of cells shows that output fatigue occurs on several timescales, with different underlying mechanisms, as well as layer and cell-type specificity (Sanchez-Vives et al.,

[1]Further alternatives include stimulus-specific adaptation, familiarity effect, adaptive filtering, mnemonic filtering (Ringo, 1996).

2000b; Ahmed et al., 1998; Sanchez-Vives et al., 2000a). For some paradigms, in particular for prolonged stimulation, this may be the underlying mechanism of repetition effects. For other paradigms, there is substantial evidence that output fatigue cannot be the reason for repetition effects, since they are stimulus specific. In early areas, a particular cell can show response reductions to one stimulus, but also response increases to another (Solomon and Kohn, 2014). In higher order visual areas, two different images that lead to the same initial response magnitude will not cross-adapt, i.e. one image decreases responses to a following presentation of the same image but not the other (Sawamura et al., 2006). Similarly, optogenetic stimulation of a cell in IT has little or no effect on image repetition effects (Fabbrini et al., 2019). Importantly, showing stimuli repeatedly that do not evoke responses in an IT cell can increase its responses to other stimuli, pointing to repetition effects as a network process (Kaliukhovich and Vogels, 2016).

Input fatigue and synaptic depression. A related but distinct notion is that neurons in a particular area respond less due to fatigue in their inputs (Vogels, 2016; De Baene and Vogels, 2010). This includes inputs from upstream areas, i.e. an inheritance of an effect generated earlier in the system. These may show output fatigue. However, subcortical inputs tend to show substantially less repetition effects than cortical areas (Solomon and Kohn, 2014; Sanchez-Vives et al., 2000b). In higher areas, a certain degree of position invariance of the effect implies a local origin, since lower areas have smaller receptive fields (De Baene and Vogels, 2010). It is also possible that very small changes due to local output fatigue can be amplified in the network (Vogels, 2016).

In addition, the synapses between the inputs and the targeted neurons might depress. This can be equally true for inputs deriving from the local network. Synaptic depression is variable between cells and exists in several forms with varying timescales, although they are typically in the sub-second range (reviewed in Fioravante and Regehr 2011). Notably, this particular form of "fatigue" may have strong functional relevance. Synaptic depression has been suggested to be a crucial form of gain control. It enables a cell to detect changes in the input that are synchronous in time, regardless of the input rates of different afferents (Abbott, 1997). This results in strong transient responses to synchronous input changes, even if the rates of the afferents are so slow that they would normally be masked by noise in more active afferents. Functionally, this translates to a sensitivity to changes in the visual stimulus, possibly at the cost of some precision of the encoding. Additionally, once synaptic depression is established, neurons may be able to compute multiplications, rather than just additions of inputs (Rothman et al., 2009). Synaptic depression therefore qualitatively changes the computational abilities of individual cells. Furthermore, it can facilitate synchronization between cells (see below). Currently, little is known about synaptic depression in awake animals. Its weakening during the "up-state" in anesthesia (Reig et al., 2006) and its dependence on neuromodulation (Gil et al., 1997) suggests that its role in awake animals could be less (or more) significant than expected based on slice-based and anesthetized studies.

Input or output "fatigue" proposals concern themselves most directly with circuit mechanisms (the implementational level), with any effect(s) on function seen as the result of such a low-level mechanism. Proposals that entail some benefit of stimulus repetition include the concepts of "sharpening", "facilitation", "prediction", "change detection" and "synchronization" (Gotts et al., 2012; Grill-Spector et al., 2006; Solomon and Kohn, 2014).

Sharpening. Sharpening refers to a selective reduction in responses that are least

informative to downstream areas, resulting in more efficient responses, possibly with an improved signal-to-noise ratio (Desimone, 1996; Wiggs and Martin, 1998; Grill-Spector et al., 2006). In a rate-based sharpening scenario, neurons that are weakly responsive to begin with should show stronger response reductions. There is substantial evidence for sharpening in the case of prolonged training with particular stimuli in the higher levels of the hierarchy (Rainer and Miller, 2000; Baker et al., 2002; Freedman et al., 2006). In case of repetitions within a single day, a study that set out to explicitly test the sharpening model found no evidence for this to occur in IT (De Baene and Vogels, 2010). Rather, several studies have reported a "scaling" effect, with larger response decreases when cells are more strongly driven, such that relative responses are maintained (McMahon and Olson, 2007; De Baene and Vogels, 2010; Miller et al., 1993).

Facilitation. Facilitation refers to a suggested mechanism that accelerates stimulus processing, such that overall neuronal responses are decreased (James et al., 2000; James and Gauthier, 2006; Henson, 2003). This proposal includes reductions in response latency, and in response duration. Unlike the output fatigue and sharpening models, there is no explicit prediction regarding any relationship between a neuron's responsiveness and later change. To date, decreases in response latencies have not been observed with stimulus repetition in single neuron responses (reviewed in Gotts et al. 2012). However, response reductions with stimulus repetition can have a characteristic timecourse, with the largest reduction a few hundred milliseconds after the reponse onset (Vogels, 2016). Arguably, this could be interpreted in line with a facilitation model, in particular the idea that the "resolution" time is reduced until the activity reaches a stable state (Henson, 2012). In higher order visual cortex, object identification may be slower and more reliant on recurrency for difficult images (Kar et al., 2019). Interestingly, in a paradigm contrasting familiar to novel images, Meyer et al. (2014) observed sharper temporal dynamics for familiar images, including a "truncation" of the response (see also Peissig et al. 2007 and Manahova et al. 2019). In summary, there is currently some evidence pointing to a potential role for faster processing with stimulus repetition based on the timing of response reductions or changes in the temporal profile of responses. A direct link to a faster resolution has yet to be established.

Predictive coding. Prediction proposals constitute a variant of facilitation. According to a particular predictive coding framework, activity in lower areas reflects prediction errors that are sent to higher areas, which generate predictions that can suppress responses in lower areas (Friston, 2005; Henson, 2003; Grill-Spector et al., 2006). With repetition, predictions could become more successful, reducing error and therefore activity in lower areas (Auksztulewicz and Friston, 2016). The predictive coding framework therefore suggests that response reductions in a particular area should involve interactions with higher-order areas of a suppressive nature. Functional interactions between higher and lower areas do appear to change with stimulus repetition. However, it is unclear that these are of the nature or content suggested by the hierarchical predictive coding framework (Ghuman et al., 2008; Chao et al., 2018; von Stein and Sarnthein, 2000; Brunet et al., 2014).

Change detection. The above proposals suggest some benefit of stimulus repetition for the encoding of the repeated stimulus. Alternatively, stimulus repetition may benefit other stimuli, or in other words, change detection - possibly even at the expense of precise encoding of adapted or repeated stimuli (Wissig and Kohn, 2012; Vogels, 2016; Solomon and Kohn, 2014). This is in line with stronger responses for stimuli other than the repeated stimulus, and increased disciminability between stim-

uli (Solomon and Kohn, 2014; Hansen and Dragoi, 2011; Wang et al., 2011; Dragoi et al., 2002). Wissig and Kohn (2012) suggest that surround modulation and adaptation have the same purpose: highlighting saliency. Similar suggestions have been made by Hosoya et al. 2005, Sharpee et al. 2006 and Schwartz et al. 2007, relating both adaptation and surround modulation to the efficient, predictive (not necessarily in the Fristonian sense) encoding of a changing environment. Change detection is therefore a particular function of circuits that implement efficient coding. Note that this proposal is compatible with input fatigue or synaptic depression accounts on the more mechanistic level.

Synchronization. Finally, synchronization proposals suggest that increased coordination of neuronal responses can compensate for reductions in response amplitude, leading to more efficient processing (Gotts et al., 2012; Gotts, 2003; Gilbert et al., 2010). Notably, this proposal resides on a more implementational level, but implies some functional benefit. It is for example possible that a facilitation or sparsening is brought about through synchronization. Changes in functional interactions between higher and lower areas with repetition are often assessed with frequency-specific synchronization measures (Ghuman et al., 2008; Chao et al., 2018; von Stein and Sarnthein, 2000). They are therefore in line with prediction as well as synchronization proposals. Moreover, synchronization could both counterbalance and be brought about by input fatigue. Several neural network models incorporating synaptic depression exhibit rate reductions with simultaneous increases in synchrony (Gotts, 2003; Wang et al., 2011). Similar behavior can also result from models incorporating facilitation of inhibitory, or both inhibitory and excitatory synapses (Bazhenov et al., 2005). While synaptic facilitation or depression in these models are short-term plasticity mechanisms, different mechanisms like long-term potentiation or depression (Andersen et al., 2017; Malenka and Bear, 2004) would result in longer-lasting synaptic effects. To summarise, synchronization proposals are compatible with other proposals such as facilitation, sparsening, prediction, or input fatigue. Further evidence for the synchronization proposal is reviewed below, together with an introduction of different repetition designs that will aid in structuring the available literature.

2.2.2 Overview of repetition designs and relationship to gamma-band synchronization

Stimulus repetition effects have a long history in sensory neuroscience and are studied with a variety of paradigms. Most studies focus on effects on either firing rates or fMRI signal, and are reviewed for example in Solomon and Kohn (2014), Vogels (2016) and Grill-Spector et al. (2006). Given the variety of paradigms, there is a corresponding variety of effects on gamma-band synchronization, which will be briefly reviewed here (see Table 2.1).

Classical adaptation paradigms measure perceptual or neuronal responses before and after the prolonged presentation of a stimulus (typically on the order of tens of seconds to minutes, Kohn 2007; Solomon and Kohn 2014). Prolonged presentation can lead to response fatigue in V1 (Sanchez-Vives et al., 2000b), and also in its inputs (Zaidi et al., 2012). The do-it-yourself afterimage exercise described in the General Introduction (section 1.4.1) is an example of such adaptation. Jia and colleagues (2011) showed that prolonged presentation can reduce gamma-band responses to the identical grating stimulus, and increase responses to other grating stimuli under anesthesia. Similarly, we have shown in awake animals that continous presentation of a colored background strongly reduces gamma-band responses to this stimulus, but in-

creases responses to color-opponent stimuli (see Chapter 5). Prolonged presentation paradigms can help to elucidate mechanisms of perception, using response fatigue as a tool. However, they constitute an extreme and rather unnatural form of exposure, and are difficult to study with awake animals.

Paradigm type	Name of effect or paradigm	Typical area or mode of study	Effect on GBA
Within session: Prolonged presentation	Adaptation	Retina, LGN, V1, MT; primate electrophys., psychophysics	Decrease: Jia et al. (2011), Peter et al. (2019)
Within session: one-back or within-trial immediate, single repetition	Repetition suppression, adaptation	IT, early sensory, electrophys. and fMRI MEG/EEG	decrease: IT high gamma, Kaliukhovich and Vogels (2011), increase: V1/V4, Hansen and Dragoi (2011), Wang et al. (2011)
Within session: within-trial repetition with intervening stimuli	Working memory, short-term memory	Prefrontal, IT, fMRI MEG/EEG and electrophys.	bursts in delay period Lundqvist et al. (2016)
Within session: between-trial single repetitions with intervening stimuli	Priming	fMRI MEG/EEG and rarely electrophys.	decrease: Friese et al. (2012)
Within-session: cumulative repetitions	Short-term memory, adaptation	Rare. Within-session learning effects in sensory areas	increase: V1/V4 Brunet et al. (2014), rodent/insect olfaction Laurent et al. (2001) van Wingerden et al. (2010)
Between session/long-term	Learning, SRP	IT, more rarely early sensory in primates. V1 in mouse electrophysiology.	precedes plasticity Galuske et al. (2019)

Table 2.1 — Overview of repetition paradigms. LGN = lateral geniculate nucleus, MT= middle temporal visual area (V5). SRP = stimulus-response potentiation. GBA = gamma-band activity. electrophys. = electrophysiology.

Visual stimulus repetition in awake animals is often studied on the sub-second timescale, with a lot of studies focusing on immediate and single repetitions within a trial (Vogels, 2016; Hansen and Dragoi, 2011; Wang et al., 2011; Nikolić et al., 2009). Rate decreases with repetition are widely reported for these paradigms across areas and layers (Vogels, 2016; Solomon and Kohn, 2014; Kohn, 2007; Gutnisky and Dragoi, 2008; Dragoi et al., 2002; Müller et al., 1999). In IT, high-frequency LFP signals also reduce in amplitude (Kaliukhovich and Vogels, 2011; Kaliukhovich et al., 2013; Kaliukhovich and Vogels, 2012), although it should be noted that the broad-band nature of the responses means they are likely of a "high-gamma" nature (see General Introduction, section 1.4.3). This contrasts with findings in V1 and V4 (mid-level visual cortex). Wang et al. (2011) and Hansen and Dragoi (2011) studied the responses to a briefly presented grating stimulus, when immediately preceded by a random dot pattern or another grating stimulus within the same trial, in awake monkey V1 or V4 respectively. Somewhat unusually, responses were not compared be-

tween the initial and later presentation of a grating, but only between gratings with a different history earlier in the trial. Both in V1 and V4, a grating preceded by a grating could show increased spike-field coherence compared to a grating preceded by a random dot pattern. In V1, effects were not specific to the gamma-band. Notably, coherence increases only occured when the repeated grating stimulus was similar to the adapter, which was always chosen to be similar to the preferred LFP stimulus. Stimulus specificity was also reported in V4. The average spike-field coherence and LFP power spectra for gratings did not exhibit a peak in the gamma-band. However, when nearby electrodes with similar stimulus tuning were chosen, a selective increase in the gamma-band was shown for the repeated grating stimulus. Interestingly, only cells that decreased their firing rate increased their gamma-band locking. In both V1 and V4, increases in gamma-band activity were positively correlated with neuronal stimulus discriminability. This contrasts with a lack of support for sharpening proposals for similar paradigms from IT (De Baene and Vogels, 2010; Kaliukhovich et al., 2013). Interestingly, discriminability in early visual areas improved specifically for highly similar stimuli, an aspect that is harder to quantify for the more naturalistic stimuli used in IT. The present evidence provides a tentative link between gamma-band synchrony and discriminibality, in line with change detection and sharpening proposals of stimulus repetition. This link may be specific for fine discrimination (see also General Discussion).

In within-trial, rapid repetition designs, specific measures may be taken to attempt to reset any cumulative repetition effects through intervening stimuli, masks and breaks (Vogels, 2016). Cumulative effects of repetition, and memory effects past the within-trial repetition, are expressly not of interest in such designs. Studies that *are* concerned with some form of memory of stimulus repetition across trials or interfering stimuli fall in the categories of 1) priming, 2) working memory studies, 3) long-term (cross-day) learning, and 4) studies of short-term memory or within-session cumulative response changes. Cases 1-3 will be considered relatively briefly, and case 4) will be treated more elaborately, due to the similarity with the design in the current study.

Priming refers to behavioral advantages or changes after a second presentation of a stimulus. A link between priming and repetition suppression (i.e. response reduction) has been suggested repeatedly (Wiggs and Martin, 1998; Henson, 2003; Grill-Spector et al., 2006; Gotts, 2003; Gotts et al., 2012). A paradigm that is favored in human studies of priming uses several, sometimes many, intervening stimuli between isolated (and thereby mostly just single) stimulus repetitions. Higher-order visual areas in non-human primates have been reported to show memory in the form of decreased rate responses even on these time-scales, which match human abilities to remember isolated encounters of visual stimuli with a lot of "interference" of other stimuli (Meyer and Rust, 2018; Li et al., 1993; McMahon and Olson, 2007). Across a larger population, such stimulus-specific response decreases could be used to determine the familiarity of a stimulus (Meyer and Rust, 2018; Sugase-Miyamoto et al., 2008). Reports of behavioral priming are very rare in non-human primates. In a remarkable case where a behavioral effect was found and simultaneous IT recordings were avaliable, no relationship to repetition suppression was observed (McMahon and Olson, 2007). Using MEG in humans, Friese et al. (2012) reported a decrease in gamma-band activity in visual areas for a single repetition of a line drawing for a paradigm with 2-3 intervening stimuli. No direct correlation to behavioral priming was tested.

Priming studies typically make no statement as to how a memory is maintained. In contrast to priming studies, studies that attempt to do so typically study working memory processes. In one typical variant of a working memory task, a stimulus has

to be matched to an identical target stimulus shown earlier in the trial, and different stimuli are shown in between. Such a task involves an active, within-trial maintenance over some seconds. Such a process can have functional signatures, such as elevated firing rates in higher-order brain regions (Constantinidis et al., 2018; Lundqvist et al., 2018; Fuster and Alexander, 1971; Goldman-Rakic, 1995) or so-called gamma-band bursts (Lundqvist et al., 2016), but may also rely on synaptic plasticity (Mongillo et al., 2008). One of the most influential early studies on working memory reported response reductions for repeated stimuli under such conditions (Miller et al., 1993). The term "repetition suppression" was coined in the wake of these discoveries (Desimone, 1996). Notably, when the stimulus in memory was not the only stimulus that could repeat, firing response increases were found upon target repetition, whereas other repeating stimuli still showed response reductions (Miller and Desimone, 1994). This illustrates that the effects of stimulus repetition can depend on task demands. Working memory is a complex process involving interareal coordination and various oscillatory processes (Miller et al., 2018) that is historically linked to stimulus repetition, but developed into its own field of study.

On the other extreme, long-term learning is typically thought to involve many repetitions across many days, and require sleep (Schoups et al., 2001; Meyer et al., 2014; Woloszyn and Sheinberg, 2012; Tononi, 2009; Huang et al., 2018). In primates, decreased responses with increasing long-term familiarity have been reported primarily in IT, but also in prefrontal cortex and V2, an area directly downstream of V1 (Huang et al., 2018; Meyer et al., 2014; Woloszyn and Sheinberg, 2012; Rainer and Miller, 2000). Effects of familiarity on synchronization are largely unclear. Recently, gamma-band activity was linked to the reshaping of orientation preference maps in V1 (Galuske et al., 2019). In mice, massed exposure to high-contrast gratings *increases* responses in V1 over days (Cooke and Bear, 2010, 2014; Kaplan et al., 2016; Cooke et al., 2015). This stimulus-specific response increase requires NMDA receptor[2] activity. Again, its effects on synchronization are unclear. A different group reported increases in low-frequency synchronization in V1 in response to a stimulus that had been repeatedly paired with another stimulus in mice (Kissinger et al., 2018).

Finally, there are paradigms that consider short-term and cumulative forms of plasticity within a session. In IT, responses to stimuli show a continuous decrease also in such a case (Li et al., 1993; Sawamura et al., 2006). Sawamura et al. (2006) showed that for continuous repetitions of an image with maximally one other image in between, responses decreased strongly for the first few repetitions in a trial and less strongly thereafter. Changes in synchronization were not studied here, though very similar paradigms with just one repetition reported decreases also in high-frequency LFP responses (Kaliukhovich and Vogels, 2012; Kaliukhovich et al., 2013).

Recently, Wang and Dragoi (2015) reported that in mid-level area V4, during the rapid orientation discrimination learning of natural images, low-frequency (theta) spike-field coordination (spike-field coherence, SFC) increased as discrimination thresholds decreased. This effect was specific to stimuli that were not fully familiar, such that discrimination learning was still taking place. Wang and Dragoi (2015) suggest that this increase in theta SFC may be linked to memory formation or improved interareal coordination during the learning phase, rather than the optimization of local responses. Gamma-band synchronization was not reported to change with this task. Of note, the SFC analysis was less fine-grained than in Wang et al. (2011), where the

[2]N-methyl-D-aspartate receptor, a ionotropic glutamate receptor associated with plasticity (Malenka and Bear, 2004).

same group suggested that gamma-band SFC in V4 is restricted to cells with similar stimulus preferences.

In V1, Lazar et al. (2018) demonstrated that repeated stimulus exposure can lead to changes in both response strength and stimulus discriminability. Briefly flashed, high-contrast images of letters tended to induce increases in responses over the course of a session. At the same time, response variability decreased, and discriminability both on the single-cell and population level was improved. Interestingly, these changes were confined to the later part of the response, rather than the initial response transient. Furthermore, this study demonstrates that adaptive cumulative changes can occur in V1 even under anesthesia, i.e. in the absence of changes in attention or intentional learning.

Brunet et al. (2014) were the first to report cumulative gamma-band increases in early and mid-level visual areas with stimulus repetition. Coordination between V4 and V1, V4 SFC as well as LFP power in both V1 and V4 increased with the massed repetition of grating stimuli during an attentional change detection task. In V4, where spiking activity was available, rates decreased while spike-field coordination in the gamma-band increased. Putative inhibitory interneurons in particular increased gamma synchronization, and in contrast to putative excitatory neurons, did not show a detectable decrease in rate responses. Furthermore, whilst the spike-field coordination of excitatory neurons dropped on average, the most driven cells showed this behavior the least. In other words, weakly driven cells selectively dropped out of the rhythmic activity. Responses were therefore more selectively gamma-synchronized with repetition. The repetition effect built up over the course of the session in a log-linear manner, such that initial repetitions resulted in the strongest response increases. At the same time, the first 10 or so repetitions at the beginning of a session showed a relatively weak increase. Since repetition number and session time were not dissociated, it is possible that the animals showed a somewhat different behavior for the first few trials in a session (for example a higher error rate), such that it is difficult to interpret these first few trials (see also Discussion).

In insect and rodent olfaction, repeated stimulation has been reported to increase high-frequency activity. In rodents, this activity was reported in orbitofrontal cortex during a learning paradigm and correlated with behavioral performance (van Wingerden et al., 2010). Odor-related gamma-band activity increased over the course of the ca. 20 repetitions tested. Since attention is known to modulate gamma-band activity, it is possible that this reflects a form of learning to attend. However, it could also be related to the stimulus repetition effects observed in primates, where attentional conditions were stable and stimuli had been overtrained (Brunet et al., 2014). In insects, activity peaking at about 20 Hz increases with repeated odors (Laurent and Davidowitz, 1994; Laurent et al., 1999, 2001; Stopfer and Laurent, 1999; Stopfer et al., 1997; Bazhenov et al., 2005). Spike-field coordination increases while spiking decreases simultaneously in a stimulus-specific manner. Interestingly, spike rates drop more strongly for the initial than for later presentations, whereas LFP power linearly increases and then asymptotes. Effects transfer between similar odors ("spill-over", Stopfer et al. 1997). Moreover, effects are independent of stimulus duration (between a few hundred milliseconds to several seconds). The effect survives intertrial durations of tens of seconds and lasts for about (has memory for) half a minute after 10 repetitions (Stopfer et al., 1997). Due to the differences in species and sensory system, and lack of knowledge of underlying mechanisms, it is unclear if these effects are related. The prima facie similarity suggests the possibility of a preservation across species and systems, which would indicate a fundamental mechanism.

The present chapter will investigate the effects of stimulus repetition in V1, using a broader set of stimuli than Brunet et al. (2014). Stimulus specificity of the effect will be tested with stimuli initially novel to the animals. The number of repetitions will be closer to that in previous insect and rodent studies, and dissociated from the beginning of the session. In the following chapter, the effect of location and the persistence of the effect (memory) will be investigated.

2.3 Methods

Since this study represents a collaborative effort of several labs, recording equipment, implanted recording devices as well as task software differed somewhat between animals.

2.3.1 Animals

All procedures complied with the German and European law for the protection of animals and were approved by the regional authority (Regierungspräsidium Darmstadt). All animals were group-housed in enriched environments with access to an outdoor space and continued to live in the facility post recording in their groups. Animal welfare was monitored by veterinarians, technicians and scientists throughout the study. Four male rhesus monkeys weighing between 12-16 kg aged 9-11 years were recorded from for this task for 2-4 weeks. All monkeys had previous experience with other forms of visual fixation and change detection tasks. Training for this particular task was minimal (if any) and did not involve the stimuli analyzed here. All recordings were from the foveal and parafoveal (up to about 8 dva eccentricity) regions of V1 (and dorsal V4 in monkeys T, H and K) and were made using chronically implanted devices (see below).

2.3.2 Surgical procedures and implants

All surgeries were performed under general anaesthesia using standard techniques according to authorized guidelines. This included peri-surgical analgesia and monitoring. All animals received a titanium implant to immobilize the head during recordings. Monkeys H and K also were implanted a recording chamber in addition to Blackrock multielectrode arrays and a connector for these arrays (Blackrock Microsystems, Hannover, Germany, see also Peter et al. 2019). Data recorded for this experiment are from 64-channel arrays (Blackrock Microsystems, inter-electrode-distance 400 μm, tip size 3-5 μm, impedance 70-800 kOhm at 1000 Hz). A reference wire was inserted under the dura towards parietal cortex. monkey A was implanted with a semi-chronic microelectrode array, with 32 electrodes (Gray Matter Research SC32-1, inter-electrode-distance 1.4 mm, impedance 0.5-2 megaOhm). The microdrive chamber was used as the reference. monkey T was implanted subdurally with a micromachined electocorticographic electrode array (ECoG, IMTEK & BCF, University of Freiburg, Germany, Rubehn et al. 2009) covering a large part of the lower right quadrant of V1 and dorsal V4. The 252-channel ECoG array consisted of 10 μm thick polyimide foils with 0.3 μm thick platinum electrodes in a polyimide grid. Global references were positioned over V1 and V4 in the same hemisphere.

2.3.3 Neurophysiological recording setup

Monkeys performed a change detection task (see section Task) seated in a custom-made primate chair in a darkened and (except monkey A) sound-attenuated booth. The animals were positioned 64-80 cm in front of a 22 inch 120 Hz LCD monitor (Samsung 2233RZ, Wang 2011; Ghodrati et al. 2015), backlit by a cold cathode fluorescent lamp. Stimulus onsets and changes were recorded with a custom-made photodiode. Eye movements from one or two eyes and pupil size (except monkey A) were recorded using infrared illumination. To this end, eye signals were calibrated before each recording session (in case of monkey A more rarely) using a standardized fixation task. Eye data was recorded with an Eyelink 1000 system (sampling rate 1000 Hz for monkey T, 500 Hz for monkey H and K) or a Thomas Recording system (ET49-B, 122 Hz, monkey A). In the case of monkey A, lever responses were additionally recorded. The animals head was restrained to allow both eye tracking and stable neurophysiological recordings. LFP and (where possible) multiunit data was collected using chronic or semi-chronic implants. Correctly performed trials were rewarded with a drop of diluted fruit juice delivered with a solenoid valve system. Recording sessions lasted ca. 1-2.5 hours depending on when the monkey stopped to initiate trials.

2.3.4 Task

The task structure reflects a compromise between aiming for more naturalistic conditions with the concomitant broader external validity and experimental control. Briefly, monkeys performed a change detection task on isolated naturalistic images, with repetitions of the images occurring only between trials, pseudorandomly interleaved with presentations of other repeating images.

Stimuli

Stimuli were taken from Hemera Photo-Objects Vols. 1, 2, and 3 (Hemera Technologies, see also Woloszyn and Sheinberg 2012). These are high-resolution color images of isolated objects (see Figure 2.3, no background, i.e. the background will be the chosen monitor color). Stimuli were presented unaltered, i.e. without attempting to equate the images basic statistics such as the average brightness. We refrained from such control since for colored natural images, such manipulations resulted in highly unnatural coloring that would impede image recognition. Preselecting images with similar brightness (as an example) on the other hand would necessarily strongly limit the color range that could be used, which would reduce their distinctiveness and thereby reduce our ability to test for stimulus-specificity of our effects. Images were selected to be perceptually distinct from one another for human observers, but no image metrics were computed to ensure this. However, the eye movement behavior of the animals strongly indicates that they were capable of both distinguishing and remembering the different images (see Results, Figure 2.4). The number of images, 25, was chosen such that for the given number of repetitions (20 per image), it was highly likely that the animals would complete at least one set on a recording day (500 correct trials per set), so that data could then be compared across days. Given this compromise between number of images, number of repetitions and likely performance, it was impossible to select an image set of a sufficient size to be representative of the full range of natural visual input. The stimulus set contained fruit, vegetables, leaves,

flower blossoms and sweets whose predominant colors were either red, orange, yellow, green or "dark" (brown, black, dark blues and greens). Note that the category of sweets was likely least familiar to the animals and is not necessarily recognized as edible (although color is a strong predictor of edibility for primates for novel items, Santos et al. 2001) There were 25 objects in total, 5 in each category and color in a 5x5 design (with the exception of monkey A, where some stimuli of the original set were replaced with images that had more inner structure, see Figure 2.3). The chosen colors reflected the dominant colors available for the categories in the stimulus set, and also ensured that there were both stimuli darker and brighter than the background. Stimuli were positioned such that they typically overlapped (slightly) with the fixation point and therefore the fovea (Figure 2.1). The center of gravity of the images was in the lower right quadrant, in accordance with the recorded receptive fields. Stimuli overlapped in a region of about 8 dva diameter. Stimulus size and position were such that all available receptive field positions in both V1 and V4 (where available) were stimulated. This means that the stimuli were not optimized for a particular RF size or location and therefore typically encompassed both the RF and its surround (see Figure 2.1), as would frequently happen during natural vision. Overlap with the fovea was chosen to stimulate the foveal receptive fields that were available in monkey T. Further, we were under the impression that the animals were better able to maintain central fixation if they were thus enabled to foveate the object to be monitored. Lastly, a foveated object may naturally attract attention (Cheung et al., 2016; Dowdall et al., 2019), and given cortical magnification of the fovea (see General Introduction, Figure 1.2) engage large parts of the ventral visual stream. Since the relative position of V1 and V4 receptive fields, as well as the average eccentricity of V1 receptive fields and their dispersion, varied between animals/recording techniques, we did not attempt to provide identical stimulation to certain receptive fields between monkeys. As a result, somewhat different parts of a stimulus will have stimulated receptive fields in each area per monkey.

Single-trial structure

On a given trial, a single image of a natural object (see Stimuli) was presented after a baseline fixation period of 1.3-1.4 s for a duration of 1.5-3 s, after which a small local change occurred (see Figue 2.2; see Table 2.2 for parameters that had some variation between animals). The monkeys responded to this unpredictable change either by lever release (monkey A) or a saccade to the change location (all others). Correct responses were followed by a juice reward during the presentation of a grey background screen. The monkeys initialized trials by acquiring fixation, resulting in some variability in the inter-trial interval. The variable stimulus time interval followed a Weibull distribution ($f(t|a,b) = \frac{b}{a}(\frac{t}{a})^{b-1}e^{-}(t/a)^{b}$, a = 0.27, b = 2, all t > 0) with a peak probability of a change occurring after around 2 s (for variability between animals, see Table 2.2. This results in a linear increase in the hazard rate (chance of a change occurring at any given moment) compared to the typical exponential increase with a uniform distribution. The linearization was chosen to reduce known hazard rate effects in gamma-band activity (Schoffelen et al., 2005; Lima et al., 2011). There were no so-called "catch" trials without changes, since the monkeys tended to produce few responses that could be interpreted as false alarms possibly because uncertainty in both space and time rendered a guessing strategy unsuccessful. After this variable duration, a small Gaussian contrast decrement appeared at an unpredictable location on the object. The possible change locations were constrained to positions outside the

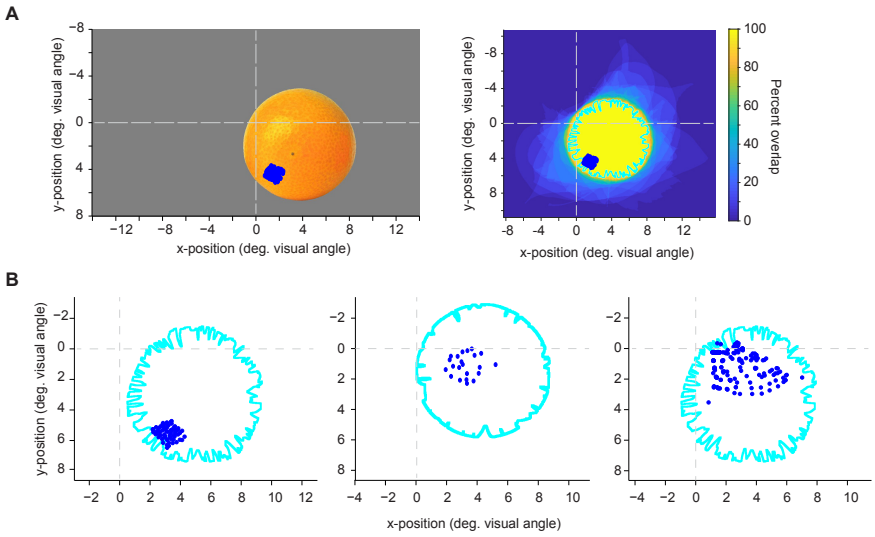

Figure 2.1 — (A) Illustration of stimulus position with respect to fixation (i.e. 0,0 refers to the fixation point) and the receptive field centers of the V1 array for an examle animal (monkey H). Left panel: Stimulus position on the monitor for an example stimulus. Right panel: Illustration of percentage of overlap between all images, the 100% region is indicated with cyan outlines. (B) Illustration of overlap of all stimuli with respect to fixation and the receptive field centers of the V1 recording sites for monkeys K, A, T (from left to right).

fixation window where all objects overlapped, which was around their center of mass in the lower right visual field and excluded the borders of the object (see Figures 2.3, 2.1). Overall task difficulty was adjusted to the performance of the individual animal rapidly on the first recording day (see Table 2.2). For one animal that was highly trained on contrast change detection for different tasks, contrast decrements were initially set for each object based on human psychophysical thresholds (monkey T, data not shown). For the other animals, maximal contrast decrements for Gaussians of a given size had to be used. Detection difficulty was therefore not equalized between images. Investigating responses to objects novel to the animals precluded adjusting difficulty per animal and object up front. This approach is in any case limited due to changes in local image structure and eccentricity within each image, both of which affect task difficulty.

Cross-trial structure

Repetitions of images were implemented across rather than within trials. Notably, this means that repetitions had no direct behavioral relevance for the animals. Specifically, on each day, the same 25 stimuli were repeated 20 times each in a pseudorandom way with a constrained lag (maximally 4 other images between one stimulus and its repetition). This was implemented such that at a given moment, one of three possible stimuli was shown (monkey A: two possible stimuli). The number of interleaved stimuli was a direct result of the constrained lag. Lag was restrained because we assumed that in order for the repetition effect to build up over trials, time between repetitions should not be too long, particularly in early visual areas (Ringach and Shapley, 2004; Kim et al., 2019; Nikolić et al., 2009). The pseudorandomly interleaved sequence of stimuli was

M.	Stimulus duration	Change difficulty	Fixation point, radius	Display software, response type
A	1.8-3 s, peak ca 1.9 s	100%, 1 dva (SD 0.2)	blue dot, ca. 1	custom software, lever release
H	1.8-3 s, peak ca 1.9 s	100%, 1.27 dva (SD 0.22)	blue dot, 1.2	MonkeyLogic v. 2013/08 on Matlab2011a, saccade
K	1.8-3 s, peak ca 1.9 s	100%, 0.82 dva (SD 0.16)	blue dot, 1.2	MonkeyLogic v. 2013/08 on Matlab2011a, saccade
T	2.3-4 s, peak ca 2.5 s	stimulus dependent	white Gaussian, 1.2	Arcade (initial version) on Matlab2014a, saccade

Table 2.2 — Task parameters that varied between animals. RT is reaction time. Fixation is in dva (degrees visual angle). M. = monkey.

generated by randomly drawing from n (2-3) possible stimuli, with the constraints that 1) a stimulus reaching maximal lag was automatically shown and 2) the likelihood of drawing the stimulus that occurred at the most recent two positions in the sequence was reduced (weighted random sample). When an image completed its 20 repetitions, another image was introduced into the group of stimuli that was drawn from. In this way, new images were introduced pseudorandomly over the course of the session. The sequence of each day started with two additional "dummy" stimuli with a randomly drawn, reduced number of repetitions. Therefore, a sequence would start with one of the stimuli of the set interleaved with the dummies. The dummies would in turn drop out before the completion of the initial stimulus of the set, so that the next stimulus of the set did not start its sequence simultaneously with the first. As a result, a given image's start in the session was dissociated from that of the other images. If the monkey completed 20 repetitions of each of the 25 images ("block"), the task seamlessly proceeded with another sequence, such that the last images in the block were interleaved with the first images of the second block (no dummies). The images chosen for the second block were produced in an identical manner, with the constraint that they were not identical to the last images in the first block. In cases where this second block was also completed, the last images were interleaved with dummy images. On the second block, a break of ca. 15 min was sometimes included at a pseudorandom time at roughly 1/3, 1/2 or 2/3 of the span of block 2. This was done in an attempt to investigate a possible reset of repetition effects as reported for a single session in one monkey in Brunet et al. (2014). Since monkeys sometimes do not reinitiate the task after a break, this was only attempted in the second block, which was not reached by all animals. The number of repetitions was chosen as a compromise between a large stimulus set and a number of trials that could be used to robustly estimate repetition effects given the many other sources of trial-to-trial variability of activity in neocortex. Errors were handled such that correct trials, misses, and responses where the target change was initially saccaded to but not held in view, were all counted as repetitions for the purpose of increasing the counter during the task. In all monkeys except monkey T, fixation breaks during the stimulus period resulted in the stimulus

turning off immediately to indicate the error. Since monkey T was well-trained on similar tasks, the stimulus was left on the screen for 1 s which allowed the animal to explore the image and which reduced overall fixation break behavior. In all monkeys except monkey T, trials where the monkey fixated correctly for > 1 s during the stimulus period were also counted as repetitions. This rule was implemented because most animals would otherwise occasionally show persevering fixation breaks toward specific images (which differed between animals). Perseveration could have destroyed the de facto lag structure between images and was often directly followed by a complete cessation of trial initiation, even for other images. Achieving 20 correct trials per block and image would have been desirable, but would probably have required training with the images. For monkey A, due to limitations in the custom presentation software in the respective laboratory and the strongest tendencies for perseverance, pseudorandom draws occurred only between two possible stimuli, and the next two stimuli were introduced after both these stimuli had completed 20 repetitions.

Cross-session structure

To dissociate stimulus-specific repetition effects from potential time-in-session effects, the order in which specific images appeared in the sequence of the day was varied across days (Figure 2.2). For example, the image of the pepper would appear at the beginning of the sequence on some days, toward the end on other days and toward the middle in other sessions. To generate pseudorandom session sequences of images (i.e. the order of appearance in the session), the first session sequence of images was randomly drawn from the set of images. The following session sequences were drawn pseudorandomly, with a reduced likelihood of positioning an image at the previously used positions in the sequence as well as their neighbors. This procedure was applied for the two blocks independently. This means that the same image occurred at different points of the session on different days, and had different neighboring images on different days. Additionally, the lag structure per image varied across days, such that the lag was independent of repetition number and image.

2.3.5 Data analysis

All analyses were done in Matlab (The MathWorks) and using FieldTrip (Oostenveld et al., 2011). All randomization or permutation tests were performed with 1000 permutations. All log-transforms have a base of 10.

Data collection and preprocessing

For monkeys H and K, recording channels were amplified, filtered between 0.05 Hz and 10 kHz and digitized at 30 kHz directly at the connector using a CerePlex E headstage (Blackrock Microsystems). Signals were then transfered out of the electrically isolated booth via optic fiber and recorded using a CerebusTM Neural Signal Processor. MUA was estimated from the broadband signal by band-pass filtering (300 Hz-30 kHz) with an 8th order zero-phase Chebyshev-filter, rectification, and low-pass filtering and downsampling to 500 Hz using the Matlab function decimate (8th order zero-phase Chebyshev-filter). LFP signals were estimated by only low-pass filtering and downsampling. For monkeys A and T, data was acquired using Tucker Davis Technologies (TDT) systems. Data were filtered between 0.35 and 7500 Hz (3 dB filter cutoffs) and digitized at 24,414.0625 Hz (TDT PZ2 preamplifier). For monkey A,

MUA was estimated by band-pass filtering (300 Hz-12000 kHz) with a 4th order zero-pass Butterworth filter, and filtering and downsampling to 1/24th of the original sampling rate using an 8th order FIR filter. Similar prelimary analyses for monkey T showed no stimulus-evoked responses in this signal, indicating that MUA could not be detected with the ECoG array at this stage of the recordings. For both monkeys A and T, LFP signals were estimated by only low-pass filtering and downsampling to 1/24th of the original sampling rate. The resulting MUA signal is a quasi-continuous measure of high-frequency field power (MUA envelope) and has been used previously by other labs (Schmid et al., 2013; Self et al., 2013; Xing et al., 2012; Legatt et al., 1980). For the calculation of rate modulations, the MUA signal was smoothed with a Gaussian kernel with an SD of 20 ms. Referencing is described in sections 2.3.2 (Surgical procedures and implants) and 2.3.5 (Electrode selection and definition of "sites").

Session selection

Monkeys A, K, and H were familiarized with a detection task on natural images using different stimuli for 1-2 days before the recordings, this data is not analysed. All animals were familiar with detection tasks before the recordings. If a monkey worked for more than 2 blocks (see Task description), any subsequent recording was excluded from the current dataset. For monkeys T, A, and H, for 1-2 recording days per animal, there were cases of either online data loss or a problem with the stimulus display software. This left the following number of recording days(blocks) per animal: monkey A 10(10), monkey K 11(20), monkey H 10(19), monkey T 10(19).

Electrode selection and definition of "sites"

To be included in analysis, channels/sites had to fulfill the following minimal criteria.

1. The site had to have a clear receptive field in the MUA activity, or in the case of the ECoG monkey, T, in the LFP responses. See section 2.3.5 (Receptive field estimation) for a description of the mapping procedures and RF quality estimates.
2. The RF had to overlap with the presented stimulus. This was ascertained a priori for all experiments by stimulus positioning. In monkey A, this criterion excluded some channels that were not positioned in foveal-parafoveal V1 (lowered below the first encounter of white matter) in an objective manner.
3. For the ECoG monkey T, some sites that showed non-physiological responses during the recordings were excluded (9/196). These sites tended to show reduced SNR in the VEP compared to neighboring sites, but due to an overall gradient in VEP strength in unipolar data, no overall threshold could be determined.

In case of bipolar derivations, we used only independent bipolar sites where both unipolar sites met the above criteria. Bipolar derivations were used for monkey T, and in explorative analyses also for the other animals (not reported here). For monkey T, local bipolar derivatives were computed between LFPs from immediately neighboring electrodes, i.e., differences sample-by-sample in the time domain, as in previous studies (Bosman et al., 2012; Bastos et al., 2015). This was done because the global references were positioned over V1 and V4 in the same hemisphere. Specifically, if

an electrode B had two direct neighbors A and C, it would only be paired with one of these to generate a bipolar site. Therefore, no unipolar recording site entered more than one bipolar site. Additionally, bipolar sites were required to both originate from the same headstage during recordings. This criterion was added because the attempt to provide an appropriate common reference across headstages failed at an unknown time point due to failures of contacts or lanes in the implant. Even though we obtained clear unipolar RFs and responses, we consider the bipolar derivation the most prudent response to this problem. This selection procedure left 90 bipolar sites in monkey T, 62 unipolar sites in both monkey H and K, and 14 unipolar sites in monkey A. The latter number is relatively low because some of the 32 electrodes were lowered into a part of V1 that covers extremely peripheral regions of the visual field. To be included in the spectral analysis for a given image, a reliable gamma peak had to be detected in addition (see below).

Trial selection

Only correctly performed trials were included in the final analyses. Trials that were part of a repetition sequence that was disrupted by a longer break (5 min) were excluded. Repetition number in a sequence could be counted in two ways: 1) counting only correctly performed trials and 2) counting all trials. If incorrectly performed trials have an influence, the second approach allocated the repetition to the "correct" position. However, since the most common type of error was a rapid fixation break, this results in missing data in the repetition sequence. We will therefore define repetition here according to the first approach. Exploratory analyses confirmed that approach (2) yields qualitatively similar results. We observed that although 20 correct repetitions per stimulus could in theory be performed, substantially less data was present after 15 repetitions according to definition 1) in many cases. We therefore restricted our analyses to the first 15 repetitions. For the regression analyses, trials that occured past a long, intentional break in the second block were also excluded, since this could potentially artificially dissociate session time from gamma-band responses through a reset (see also Chapter 3).

Behavioral analysis

Reaction times and correct versus incorrect responses were analysed using multiple linear regression analyses similar to regression analyses for neuronal data (see below). For the observation that the amount of rapid fixation breaks seemed to decrease with the first few repetitions, an exponential of the form $a + b * exp(c * iRep)$, with iRep the repetition number, was fit to the data using nonlinear least squares fitting (Matlab function fit with default settings). Confidence intervals were computed as $C = b \pm t * sqrt(S)$, where b is the coefficient, t is a threshold based on the confidence level (95%) and the cumulative student's t distribution, and S is the mean squared error*variance of a coefficient estimate.

Spectral analysis

LFP power. For monkey H and K, activity was rereferenced to the average across the V1 array for LFP power analyses. Unless otherwise noted, all spectra were computed from a fixed part of the trial, namely from 0.5-1.5 s post stimulus onset, or for the baseline, -1 to 0 s before stimulus onset. Note that the baseline period had a duration of 1.3-1.4 s, such that the chosen baseline period omits the first 300 ms,

avoiding potential nonlinearities in the response with fixation onset. The fixed length during the stimulus presentation prevents any potential effects of changing trial lengths with repetition. (Trial lengths were randomized, but since there can be dynamical changes within the trial, especially later in the trial, including different lengths might add unwanted variability). We excluded the first 500 ms after stimulus or fixation onset to minimize effects of transients and non-stationarities on the metrics of rhythmicity and synchronization. The baseline and stimulus periods were then cut into non-overlapping epochs. Two main types of spectral analyses were performed: 1) Analyses focusing on low-frequency effects using 500 ms epochs that were Hann-tapered, and 2) analyses focusing on gamma-band effects using 250 ms epochs, using multitaper spectral estimation with 5 tapers taken from the discrete prolate spheroidal sequence (Pesaran et al., 2018; Mitra and Pesaran, 1999), yielding ± 10 Hz smoothing. The combination of epoch length and taper meant that Hann-tapered analyses had a fundamental spectral resolution (Rayleigh frequency) of 2 Hz, and multitapered analyses of 4 Hz. Epochs were tapered as described and then Fourier transformed. For LFPs, relative power was computed as stimulus/baseline and then plotted on a log scale, or for regression analyses and correlations, log-transformed before averaging across sites.

MUA-LFP phase locking. For MUA-LFP phase locking, only electrodes selected by the procedure described above were used. In addition, for MUA-LFP pairs, we required that the electrodes were direct neighbors in the array. MUA-LFP phase locking was computed as follows. The cross-spectral density between LFP and MUA signal for each trial (cross-spectra) was computed using the same spectral estimation parameters as for the LFP power spectra described above. The cross-spectrum per trial was then normalized by its absolute values, resulting in cross-spectral phases (without amplitude information). Normalized cross-spectra were then used to compute the Pairwise Phase Consistency (PPC), using FieldTrip (Oostenveld et al., 2011). The PPC is unbiased by the trial count (Vinck et al., 2010b). For a given MUA site, the PPC values were then averaged across all the combinations with LFPs from the other selected electrodes. MUA-LFP combinations from the same electrode were excluded to avoid artifactual coherence due to bleed-in of spikes into the LFP (Ray and Maunsell, 2011; Buzsáki et al., 2012). Because of the distance between electrodes (at least 400 μm), this was not an issue for MUA-LFP combinations from different electrodes. Single-trial PPC values were computed using non-overlapping epochs. Exploratory analyses showed that computing PPC per repetition across sessions, or binning several trials from a repetition sequence, did not show clearer effects for the stimulus-specific correlation analyses (see below).

Determining gamma peak per image. Spectra were highly stimulus-specific, with strong variations in gamma-band peak frequency and spectral shape (see Figure 2.5). In addition, many spectra featured several distinct peaks for a single image and site (these were not due to averaging across different sessions or repetition numbers). These peaks showed up independently of doing a relative (to baseline) gamma power normalization or a $1/f^n$ correction (the latter procedure involves fitting an individual regression slope to the stimulation period spectrum per image, then multiplying each frequency bin by the slope estimate times the frequency of that bin). Frequently, they were also visible in MUA-LFP PPC spectra (see Figure 2.5).

For each stimulus, the largest peak in the gamma-band response was determined for relative power spectra (and with an identical procedure, for the MUA-LFP PPC during the stimulus presentation time). Given that the peak frequency varied across images, this could be used to determine a peak-centered gamma-band for each stimulus in

which to average gamma-band activity. Local maxima were determined between 20 and 190 Hz. A peak was defined as a position in the spectrum where spectral power was lower than the maximum on both sides of the peak as implemented in the Matlab function "findpeaks": First, the sign of the first derivative of the spectrum is computed. Then, local maxima are locations where the derivative of the sign is < 0, identifying the first point of the peak. The largest two peaks in the spectrum were collected for each site and stimulus. To ensure that the identified peak reliably occurred across days, we ran a one-sided permutation test (alpha = 0.05, n = 1000 permutations of all trials across all sessions) of the average relative power around the identified peak frequency (±8 Hz around the peak) against activity of ±8 Hz around 190 Hz for each channel and image. Note that the p-value was used as a threshold rather than for any inference about the population of peaks across sites and stimuli. The test against high-frequency activity rather than baseline was chosen to identify peaks that are reliably larger than any offset or spike leakage effects. For a given stimulus, the largest gamma peaks were typically similar across recording sites (91% of sites shared a common peak, see Figure 2.6A). We used the identified peak frequency per site and stimulus to group channels with a similar peak and subsequently analyse repetition-related effects in these channels and a frequency band ± 16 Hz around the peak. A relatively wide band was used for grouping to allow for the variability in single-trial, single-site gamma estimates (e.g. Lowet et al. 2017), potential repetition-related shifts in frequency (Brunet et al., 2014), the generally often broad shape we observed in the data with the selected multitapering (see e.g. Figure 2.5), the frequency resolution of the majority of the analyses and the known peak-frequency dependence on eccentricity (van Pelt and Fries, 2013). 49% ± 10%SD of sites with the largest peak exhibited a second peak within ±16 Hz of most common second largest peak (Figures 2.5, 2.6A). This reduction compared to the largest peak is not due to a bimodal distribution where some stimuli exhibit a reliable second peak across all sites and some show no sites with a reliable second peak. Rather, some stimuli do exhibit common secondary peaks across all sites, some exhibit very few sites with a common secondary peak, and many stimuli show a substantial percentage of sites that show reliable common secondary peaks. It is not unlikely that the detection of secondary peaks depends on signal-to-noise ratio (SNR, this notion is supported by correlations between strength of the primary and secondary peaks, data not shown). The occurrence of several gamma peaks has been reported previously for artificial stimuli like gratings in both humans and non-human primates (Hoogenboom et al., 2006; Murty et al., 2018; van Pelt et al., 2012). The observed additional peaks could reflect additional, independent phenomena, or indicate the presence of harmonics in the signal. Harmonics can occur when a signal is not sinusoidal in shape as assumed by Fourier analysis (Scheffer-Teixeira and Tort, 2016), as would for example be the case with a periodic signal that has a sharp edge when rising or falling. Harmonics occur at multiples of the fundamental frequency. In the present dataset, some peaks occurred at frequencies that were clearly too close together to be multiples of each other, whereas others likely constituted harmonics. Indeed, conspicuously many secondary peaks occurred at about twice the frequency of the largest peak (Figure 2.6). Interestingly, another, smaller group of secondary peaks occured at about half of the frequency of the largest peak (see also Murty et al. 2018). Preliminary NM-phase-coupling analyses showed that significant harmonics could be detected for some stimuli, after appropriate correction for inherent filtering-induced artificial coupling using trial randomization (Scheffer-Teixeira and Tort, 2016). Phase coupling depended linearly on power, and significant peaks were therefore more likely for stimuli with stronger gamma-band responses. This again points to a role for SNR,

but clearly demonstrates that at least some stimuli show harmonics in their gamma-band responses. Since the secondary peaks by definition had a lower SNR and seemed to constitute an inhomogeneous group of harmonics and, on occasion, other processes, all analyses involving peak detection will focus on the largest peak per stimulus.

Receptive field estimation

For all animals except the ECoG monkey T, receptive fields were mapped with moving bar stimuli (spanning the entire monitor). Moving bars (width 1/1/0.1 dva, speed 8/8/17 dva/s, monkey K/H/A) were presented in 8 orientations for monkeys H, K and 8 - 16 orientations for monkey A, each for 10 - 20 repetitions. MU responses were projected onto the stimulus screen, after shift-correction by the response latency that maximized the back-projected response. MU responses were then fitted by a Gaussian function. This Gaussian was used to extract the 10th percentile and the 90th percentile, and this was done separately for each movement direction. Across the 8/16 directions, this yielded 16/32 data points, which were fit with an ellipse. The center of the ellipse was taken as the RF center.

For the ECoG monkey, a red circular stimulus (maximal brightness, RGB [255,0,0], size: 1 dva radius) was presented on a gray background (RGB [128,128,128]) on a grid of positions (with 1 dva steps, i.e. approximately 50% overlap) in the lower right visual field as well as the fovea and the first 1.5 dva above the horizontal meridian (after an initial broader mapping that determined the coverage of the array). Since the ECoG yielded no clear MUA or high-frequency responses, receptive fields were assessed using average relative gamma power from 30-90 Hz (centered approximately around the peak in the spectrum, time window 0.3-4.5 s post-stimulus, power computed with the same parameters as the LFP power described above). To obtain receptive fields, for each channel and each location covered by the grid, relative power was computed by averaging all trials where the stimulus overlapped with the grid location. The receptive field maps of each channel were then normalized by the maximum value, smoothed with a Gaussian (0.25 dva size, SD 0.1 dva), and z-scored. The grid location with the maximal response was taken as the RF center.

For all monkeys, the backprojected data was thresholded with a z-score of 1, and activity above the 85th percentile was used as the RF estimate when computing stimulus attributes in the RF of a recording site.

Measuring shape of stimulus repetition effects on MUA responses

We observed that MUA responses showed a rapid decrease for the first few repetitions, followed by a lesser decrease for further repetitions. To quantify this effect, linear slopes were fit to the first 4 (“early”) and the later repetitions (“late”) for each stimulus and animal. The cutoff of 4 was chosen somewhat arbitrarily, in an attempt to do some justice to the stronger initial slope, and still allowing for a permutation test of the fitted slopes. Slopes were fit either to the session average responses per stimulus, or to each session individually and then averaged (both results reported in text). Several ways of normalizing the average responses are explored in the Results section. For significance testing, slopes were then averaged across stimuli, and the average difference between early and late slopes was computed. Both the average slopes and the difference in the strength of early vs late slope were then averaged across animals and tested against a permutation test. For each stimulus, the early or late repetitions were individually permuted randomly nPerm = 1000 times to generate a random dis-

tribution. For a two-sided test, the minimal p value obtainable is therefore 0.002. The reasoning for individually permuting early vs late repetitions was that a permutation of the full 15 repetitions would have frequently resulted in the initial presentation to be in the "late" part, yielding strong artificial negative slopes. In case of time-resolved analyses, the same procedure was applied to all time bins and a multiple comparison correction was applied with an alpha and false discovery rate of respectively 0.05 (Korn et al., 2004).

Measuring dependence of stimulus repetition effects on stimulus response strength

The degree to which a given recording site changes its response (e.g. MUA response) to a stimulus may depend on the overall response strength to the stimulus. We therefore analyzed the dependence of the change with repetition on response strength, for early and late repetitions separately. Several potential pitfalls have to be considered in this analysis. 1) The mean response strength and change with repetition can show trivial correlations due to circularity. For example, all other things being equal, a site which shows increasing response strength with repetition will also show a higher mean response. 2) In cases where the response strength is weak, changes in response strength are potentially limited by a floor effect. In any case, overall SNR is lower and estimates of relative changes in response will be noisier. To address the first problem, we linearly fit the responses for each session, recording site and stimulus for the repetitions in question (early or late), in a cross-validated manner. The fit $y = a + b * iRep$, with iRep indicating repetition number, yielded an estimate of the intercept at repetition "zero", as an estimate of response strength without any repetition-related increase. The fit also yielded a slope. When slope and intercept estimates are based on the same data, another bias occurs in finite data due to regression to the mean. We therefore performed two linear fits, each based on a different half of the repetitions. One half included every second repetition starting with the first, the other half every second repetition starting with the second, resulting in two independent estimates of slopes and intercepts. Simulation confirmed that this removed the bias. We then tested whether the intercept was predictive of the slope through Spearman's rank correlation. For each stimulus and recording site combination, we obtained the median slopes and intercepts across sessions. Across the stimulus-site combinations, we then correlated slopes with independently estimated intercepts. We performed this procedure for the two independent combinations of slopes and intercepts and averaged the two resulting correlation values.

To address the second problem, a potential trivial correlation due to a floor effect in stimulus-site combinations that were weakly responsive, we assessed the consistency of the correlation values across a median split of the data by response strength. For each animal, the stimulus-site estimates of the intercept as our estimate for response strength were partitioned into two parts with equally many data points, and the correlation was computed for each resulting half of the data. To test for statistical significance, correlations were then averaged for each half across animals, and the resulting correlation value tested against a permutation distribution (randomizing intercepts 1000 times for each data half per animal, then averaging across animals) using multiple comparison correction across halves (Korn et al., 2004).

Measuring stimulus repetition effects on gamma-band responses

The shape of the repetition effects of gamma-band responses was measured similarly to the analyses of MUA responses described above, and based on the gamma-peak aligned responses per stimulus. Since gamma-band responses could show decreases and increases with repetition, the distribution of early versus late slopes by the types "early and late slope increases", "early and late decreases", "early decrease, late increase" and " early increase, late decrease" was investigated and tested against chance using a chi square statistic. The observed distribution was also compared to MUA repetition slope distributions using multiple-comparison corrected pairwise permutation tests. Correlations between early and late slopes for gamma-band and MUA signals, and combinations thereof, were tested using permutation tests.

Measuring stimulus-specific repetition effects through normalized correlation

We reasoned that any stimulus-specific trajectory of a given feature with repetition (e.g. increase, rapid decrease followed by steady response, etc.) would be reflected in correlations between the trajectories across sessions: the trajectory of stimulus x_1 in session n_1 should correlate more strongly with the trajectory of x_1 in the other sessions than with the trajectory of x_2 to x_N in the other sessions. The trajectory can be computed for arbitrary features, such as LFP power in a specific frequency band, or MUA responses in a given time bin. By repeating the process for different frequencies or time bins, correlation spectra or time-resolved correlations are obtained. To compute correlations, we first normalized the trajectory for each session, site and stimulus using a z-score across the repetitions. This results in a trajectory that has a specific shape, but is both demeaned and scaled by the standard deviation of the distribution of repetitions of a stimulus in a session. The z-scored data was then concatenated for each session with a fixed stimulus order (i.e. irrespective of the stimulus order in the session), yielding a "fingerprint" vector of trajectories for each session (see Figure 2.13A for illustration). The normalization should remove effects that are caused by both greater means and greater variance (that typically accompanies greater mean responses) between different stimuli, in the feature of interest, for example in gamma-band activity. Explorative analyses showed that results based on demeaned responses were qualitatively similar to results based on z-scored responses. For a session to be included in this analysis, it was required that at least 50% of the "fingerprint" vector contained data (this excluded 0 - 2 session blocks per animal, which were incomplete second blocks of the day, range thereafter 82 - 91% of the session fingerprint contained data). Furthermore, for spectral analyses, it was required that at least 5 sites showed a reliable gamma-band peak for a stimulus to be included (this excluded 0 stimuli for gamma power, 4, 0, or 13 stimuli for PPC per animal). The data was averaged across sites after normalization. We then used a split-half procedure, where two random halves of the sessions were averaged repeatedly (s = 100 times) and the two average "fingerprint" vectors were correlated. These s split-half Pearson correlations were then averaged to yield a final estimate of the correlation value per animal. Subsequently, values were averaged across animals. To see whether this correlation value was stimulus specific, we used a permutation test. Specifically, for nPerm = 1000 iterations for each of the s split-halves, we computed pairwise correlations between the intact vector of one session half and the other session half with the trajectories reordered in a randomized stimulus order. Using the same procedure as for the observed data, the s split-halves were then averaged. Subsequently, the nPerm values were aver-

aged across animals. The resulting distribution of nPerm correlation values was then used for a two-sided test at alpha = 0.05. For spectra or time-resolved correlations, we used a false discovery rate based multiple comparison correction (FDR = 0.05, alpha = 0.05, Korn et al. 2004). Note that averaging values can lead to larger correlation values compared to pairwise correlations of single sessions. The correlation value reported here is between two average halves of all sessions.

Multiple regression analyses

Multiple linear regression models of the form $y = b_0 + b_1 X_1 + b_2 X_2 ... + b_N X_N$ were fit to single-trial gamma-band or MUA responses (y) using N predictors X_1 to X_N. Fits were performed using the Matlab function fitlm (or fitlme when including random effects terms). Final analyses were based on models with fixed effects for individual animals or sessions. Qualitatively similar results were obtained with random effects models. The final model is based on pooled data from all animals, models for individual animals also yielded qualitatively similar results for the effects of stimulus repetition. Full models that can include non-significant results are reported in text. Dropping these predictors did not change any effects qualitatively. Pairwise correlations between predictors were performed, and in case of predictors with high correlations, only one of the predictors was included in the model. Stimulus identity and recording session number were treated as categorical predictors per animal.

2.4 Results

2.4.1 Task and behavior

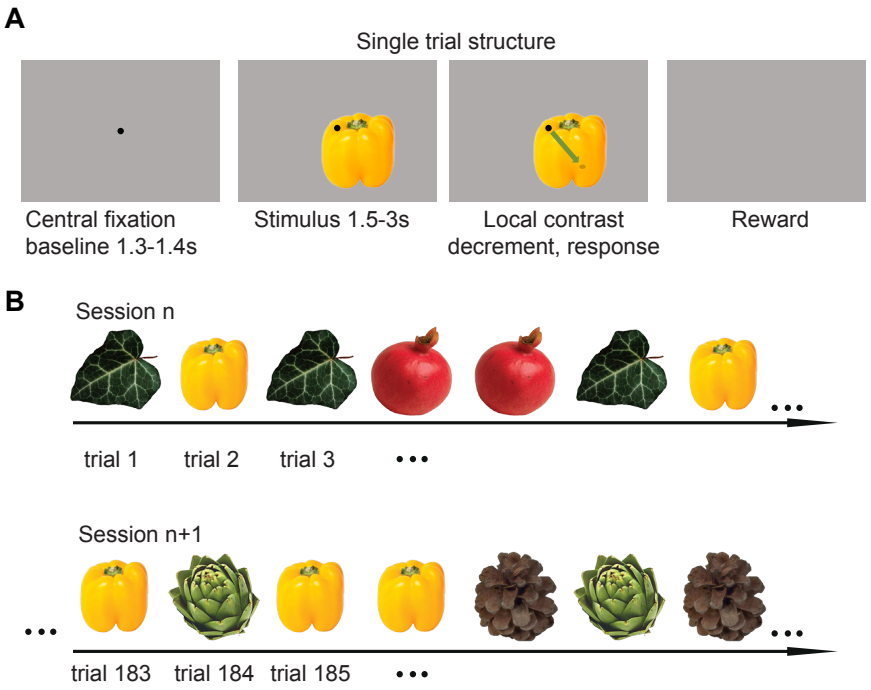

Figure 2.2 — (A) Single trial task structure illustrated with a yellow pepper stimulus. After self-initiating fixation, 1.3-1.4 s of gray background stimulation followed. Then, a single stimulus was flashed on. After 1.5-3 s (see also Table 2.1), a local contrast decrement occured at a random position on the stimulus. Monkeys responded to the change either with eye movement (monkeys T, H, K) or lever release (monkey A) and were rewarded for correct responses during a gray background screen. (B) Between-trial and session task structure. Within a session, stimuli could repeat immediately or with up to 4 intervening stimuli. A given stimulus (e.g. the yellow pepper in this example) could occur in different parts of the session on different days, and with different intervening stimuli.

We repeatedly presented colored natural images to monkeys performing a change detection task (Figure 2.2A). On each trial, one stimulus overlapping both the fixation spot and the receptive fields was presented (for 1.5-3 s, unpredictable change location and timing, see Methods, Table 2.2 for details, Figure 2.1 for receptive field locations). Repetitions occurred between trials. Stimuli could repeat immediately or with up to four intervening other stimuli (the so-called lag, Figure 2.2B, see Methods). This order was pseudo-random, effectively making it unpredictable whether the next trial would constitute a stimulus alternation or repetition. 25 different stimuli were presented 20 times each in this interleaved manner. To dissociate stimulus-specific repetition effects from 1) more general effects occurring over the course of a session, and 2) from effects arising from the precise sequence of stimulation, especially the neighboring stimuli, we ensured that 1) the position of a particular stimulus in the overall sequence varied between days and 2) the neighboring stimuli and lags between repetitions for a given stimulus varied randomly between days (Figure 2.2B, see Methods). The stimuli

were images of isolated objects, chosen to be perceptually distinct from one another in the eyes of the authors (see Methods and Figure 2.3). Stimuli could have more or less inner structure, but were always images of a single object. They roughly fell into 5 categories (leaves, flowers, sweets, fruit, vegetables) and 5 chromatic groups (dark, green, yellow, orange, and red).

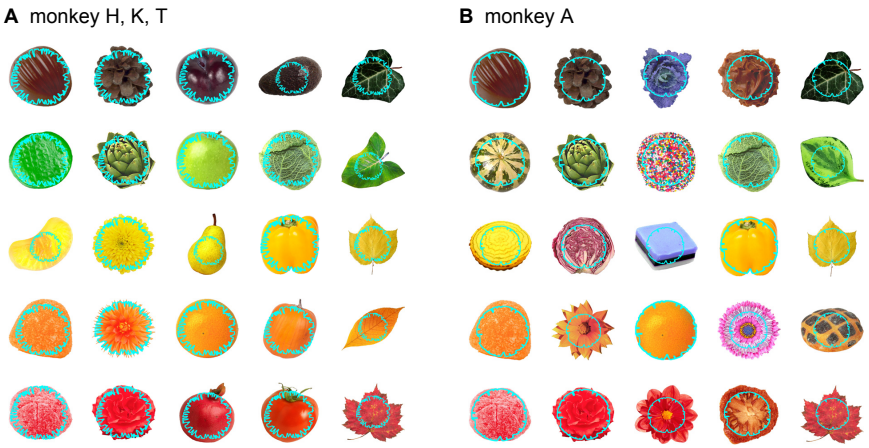

Figure 2.3 — (A) Stimuli used for monkeys H, K and T. Stimuli were selected from a large library (Hemera Photo-Objects, Hemera Technologies). (B) Stimuli for monkey A. Cyan outlines on each stimulus illustrate the overlap with all other images used for a given animal, and thereby the area where a change could occur.

The stimuli were novel to the animals on the first recording day. Although the animals were not required to memorize images and stimulus repetition was not task-relevant, there were clear indications of stimulus memory in the animals' spontaneous behavior. Specifically, the animals were more likely to respond with rapid fixation breaks during the first few presentations of a stimulus. This is indicated by a significant exponential fit to the correct trial rate (Figure 2.4A, exponential fit indicating a rapid increase in hit rate with repetition, a = 0.70, b = -0.32 , c = -0.48, r-squared = 0.95; all confidence intervals excluding zero, see Methods). By contrast, the correct trial rate was flat across repetitions when rapid fixation breaks were excluded (fit parameters not different from zero other than offset). This behavior reduced over the course of the recording days (Figure 2.4B), indicating some memory also across days. Note that due to the design of the experiment, session-novel stimuli were introduced over the entire course of the session, so this behavior cannot be explained by increased fixation breaks at the beginning of the session.

Nevertheless, a "repetition" versus a "total trial number" explanation for this behavior was also explicitly tested using a multiple regression approach. We built regression models that included monkey identity, stimulus identity and the eccentricity of the stimulus change, and as a final step included either repetition number, or total trial number in the session (see Methods). A model including repetition number outperformed models including total trial number (theoretical likelihood ratio test, final model r-squared = 0.173). By contrast, when fixation breaks were excluded, repetition number no longer had a significant effect on behavior. Taken together, these analyses indicate that the animals would transiently forego reward for attempted visual exploration. This behavior also indicates that monkeys experienced the stimuli as perceptu-

ally distinct. Based on similar regression models, reaction times showed no significant relationship to stimulus repetition, nor to total trial number (all $P > 0.3$). This could be expected given that the change location and timing was not related to stimulus repetition. Multiple regression modeling indicated "stimulus change eccentricity", "change timing" and "stimulus identity" as significant predictors of reaction times, though explained variance was low (5.4% in total, all $P < 0.002$, all r-squared < 0.03). Monkeys were slower when changes were more eccentric, and faster later in the trial on average (eccentricity, beta = 0.0032, change timing, beta = -0.01, reaction time and change time in seconds). Reaction times were demeaned for each individual session per animal. Note that the dependence on eccentricity is not an artifact of a longer time for a longer saccade, since reaction time was defined as the time point when the animal crossed the fixation boundary for all animals responding with eye movements.

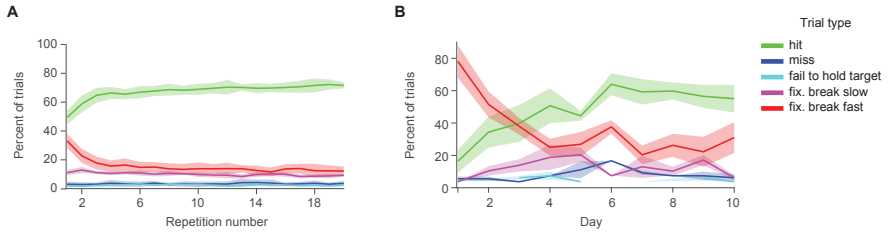

Figure 2.4 — (A) Development of behavior across stimulus repetitions. Percentage of trials that were hits, misses, cases where the change location was first acquired but not fixated for long enough, or fixation breaks that were slow (> 1 s) or fast (< 1 s). Standard error is across animals. Regression modeling indicated that there was a significant influence of repetition number on correct responses, due to increased rapid fixation breaks for initial repetitions. There was no significant influence of repetition number on reaction times. (B) as in (A), but development across days for the first presentation only. Standard error is across animals.

2.4.2 General properties of V1 responses to chromatic natural images

Before turning to repetition-related changes in responses, some understanding of the range and stimulus specificity of the typical (average) responses in V1 is required. The use of naturalistic stimuli is still relatively rare, in particular for studies of stimulus repetition (Schwartz et al., 2017), and the prevalence of gamma-band responses for such stimuli has been questioned (Hermes et al. 2014, 2019, but see Brunet et al. 2015; Brunet and Fries 2019). We therefore first describe the average V1 responses in our data, before turning to the effects of stimulus repetition. Note that due to the involvement of different labs and recording methods, LFP responses are available from four animals, MUA responses from three animals, and pupil responses from three animals.

In Figure 2.5, we illustrate the cross-site average responses to four stimuli for monkey H. MUA responses are characterized by transients of varying strength, followed by responses that can either show sustained elevated responses above the pre-stimulus responses (top row), or sustained reduction below baseline responses (middle rows). In this post-transient period (0.5-1.5 s), LFP spectra exhibited clear gamma-band responses, which however varied strongly both in overall amplitude as well as peak frequency and spectral shape. The post-transient period was chosen to minimize effects of stimulus onset transients and thereby non-stationarities on metrics of rhythmicity or synchronization. Spike-field locking spectra (pairwise phase consistency, PPC, see

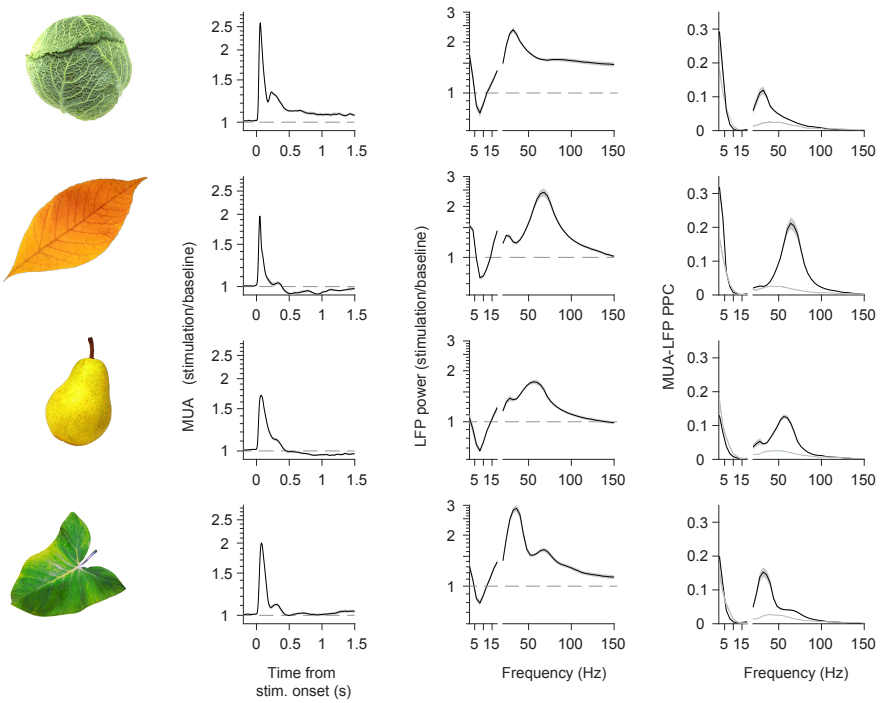

Figure 2.5 — Cross-session, cross-site average responses to four example stimuli in monkey H. Shaded area indicates standard error of the mean across sessions. Left column: MUA response (stimulation/baseline). An initial transient could be followed by sustained responses above baseline, or reductions below baseline. Central column: Average change in LFP power (stimulation/baseline). Low-frequency power spectra ($\leq$ 18 Hz) were estimated using a Hann taper based on non-overlapping 500 ms windows. High-frequency power spectra ($\geq$ 20 Hz) were estimated using a multitaper method based on non-overlapping 250 ms windows (see Methods). Right column: Average MUA-LFP locking, estimated using pairwise phase consistency (PPC). Frequency cutoffs and spectral estimation parameters as for LFP power. Gray line indicates average baseline activity for the stimulus. Note the similarity in the shape and strength of LFP power spectra and PPC spectra.

Methods, Vinck et al. 2010b, 2012) behaved similarly to the LFP spectra. Notably, we found clear PPC inspite of low MUA responses.

The averaging of spectra presupposes that the different sites are similar in overall spectral shape, including e.g. similar peak frequencies. This is indeed the case, as quantified in Figure 2.6A. On average across stimuli and animals, 91% (SD 2.9% across animals) of sites had their largest gamma peak within ±16 Hz of the most common largest peak across sites (see Methods). Conclusions were similar when grouping only sites with the peaks within ±8 Hz of the largest peak (87% ± 4.6% overlap). This justifies the use of a common gamma band of interest per stimulus. 49% ± 10% of sites that showed the largest peak also exhibited a second peak within ±16 Hz of most common second largest peak (Figures 2.5, 2.6A). This apparently included both harmonics and non-harmonic secondary peaks (see Methods for longer discussion). Peak-aligned analyses will focus on the largest gamma peak per stimulus. The variability in average peak frequency and peak gamma-band power between stimuli is shown for each animal in Figure 2.6B-C.

What are the underlying causes for this variability? A variety of stimulus attributes

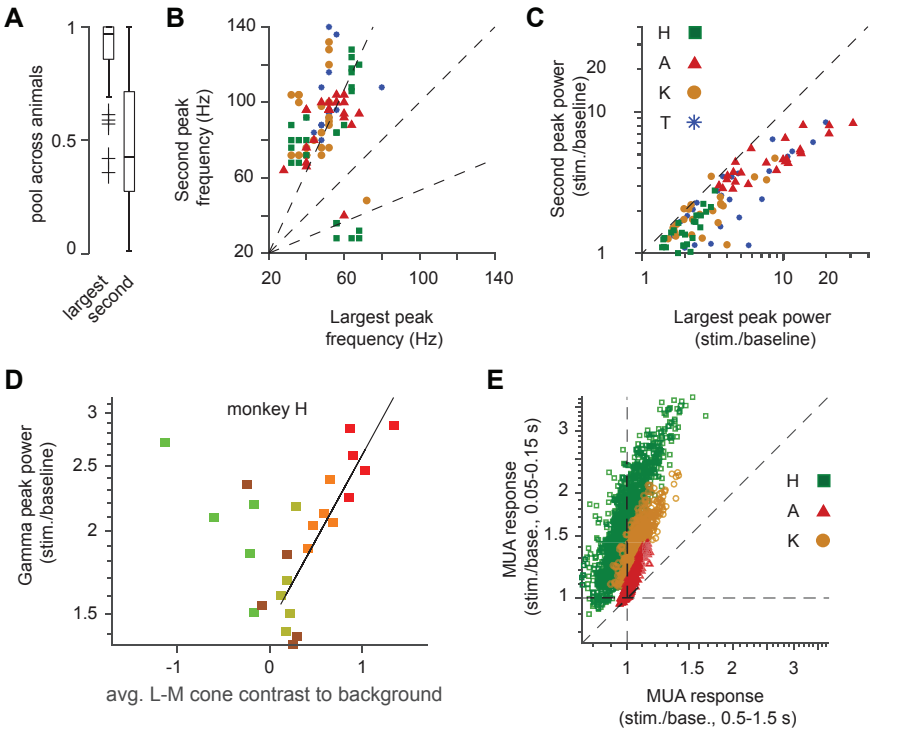

Figure 2.6 — (A) Prevalence of shared LFP gamma-band peaks (within ±16 Hz, see Methods) between sites across stimuli and animals. Largest peak refers to the largest peak most common across recording sites. Second peak refers to the most common second largest peak among recording sites exhibiting the largest peak. Horizontal lines refer to median, 25th and 75th percentile. (B) Peak frequency of the largest peak versus second largest peak. Each symbol indicates average responses for a stimulus and animal. Dashed lines indicate the ratios 1:2, 1:1 and 2:1. There is a cluster centered around 2:1, indicative of harmonics (see also Methods). However, other peaks are clearly not harmonic in nature. (C) Fold change in power at the gamma-band peak and second largest peak for all stimuli. (D) Dependence of relative gamma-band power on L-M cone contrast for an example animal (monkey H). Each symbol shows average gamma-band power around the largest peak per stimulus across sessions and sites. Color of the symbol is an approximation of the stimulus color (yellow stimuli are shown in a darker hue) to give an intuition of the color dependence. Correlation between positive L-M cone contrast and gamma power was r = 0.88 for this animal. (E) Average MUA fold change responses (stimulation/baseline) during the initial transient and responses during the post-transient period for each recording site and stimulus. Dashed lines indicate equality line and regions below baseline activity (prestimulus gray screen).

are known to influence the strength of gamma-band responses for artificial stimuli (e.g., Jia et al. 2011, 2013a; Roberts et al. 2013; Gieselmann and Thiele 2008). Some of these attributes, such as contrast, are more difficult to define for natural images (Peli, 1990). Interestingly, in natural images, contrast is strongly related to image structure, which has been recently related to gamma-band responses for grayscale natural images (Brunet and Fries, 2019). Given the small set of stimuli used here, and that no attempt was made to dissociate different stimulus attributes such as brightness, contrast and color, the question of underlying causes is difficult to pursue in this dataset. However, one effect was so prominent that it should be mentioned inspite of these caveats. As illustrated for an example animal in Figure 2.6D, positive L-M

contrast (relative redness) strongly correlated with average gamma-band responses for individual stimuli. Correlations were r = 0.85 (Pearson's r) on average across monkeys (range 0.72-0.93, all $P < 0.001$ using a t-test). As can be seen from Figure 2.6D, too few data points were present to investigate negative L-M cone contrast on this level, and the general lack of blue stimuli precluded any meaningful analysis of the other color axis. Inspired in part by the findings from this dataset, the topic of color will be investigated using more controlled stimuli in Chapters 4 and 5. Other stimulus attributes are the topic of an ongoing project using a larger stimulus set that is beyond the scope of this thesis (Uran, Peter et al., in preparation.) Figure 2.6E contrasts the transient and post-transient MUA responses for each site and stimulus, illustrating differences in the overall response strength between animals. Furthermore, it can be seen that responses were frequently reduced below baseline in the post-transient period. In summary, average responses to the stimulus set show a diverse range of responses both in terms of MUA responses and in gamma-band responses. Gamma-band peaks were detectable for all stimuli, albeit with strongly varying magnitude.

2.4.3 Stimulus repetition effects on MUA responses show a characteristic shape

We observed a repetition-related effect on MUA responses across recording sites and stimuli of a particular shape. Responses decreased most strongly for the first few repetitions, followed by more weakly decreasing responses for later repetitions. All analyses are restricted to the first 15 repetitions, since the amount of correct trials was low thereafter (see Methods). Figure 2.7A shows the site- and trial-averaged responses relative to baseline for a few stimuli of an example session. Note that the first presentation of a stimulus occurred at a random time point within a session, such that the comparatively strong decrease for the first few repetitions likely indicates that these effects are stimulus specific. Further below, stimulus specificity of the repetition effect will be tested explicitly (Figure 2.13). The initial rapid decrease followed by a later, slower decrease with repetition was consistently observed across stimuli and animals for cross-site average responses. Responses averaged across trial time were relatively low. This poses some challenges for normalizing the data appropriately and can result in floor effects for repetition effects for individual stimulus-site combinations. Several approaches to quantify the repetition effect follow here.

In the first approach, the change in response with repetition was computed for each stimulus, site and session (i.e. individual trials, average of entire stimulation period) by dividing each response (relative to baseline) by the mean across repetitions. Cross-site and cross-session averaged responses of individual stimuli for each animal are shown in Figure 2.7B. An analysis by site will follow in Figure 2.9. A value of 1.02 denotes an increase of 2 percentage points over the mean across repetitions on the average across sites for the first presentation. Note however, that response increases computed over the entire trial and site-averaged were low (ca. 10 % above baseline on the grand average, see e.g. Figure 2.7A and Figure 2.6D), such that this value is dominated by the baseline of 1. As a consequence, a change in percentage points of this scale can be relatively large when considering the change relative to the stimulus-induced response. In Figure 2.7C, responses were log-transformed to remove the strong weight by the baseline (log (stim./base.)). Additionally, responses were averaged across sites before normalization across repetitions, weighing the analysis toward more responsive sites. Furthermore, since a response ratio generates outliers when values in numerator and denominator are similar (e.g. both close to baseline),

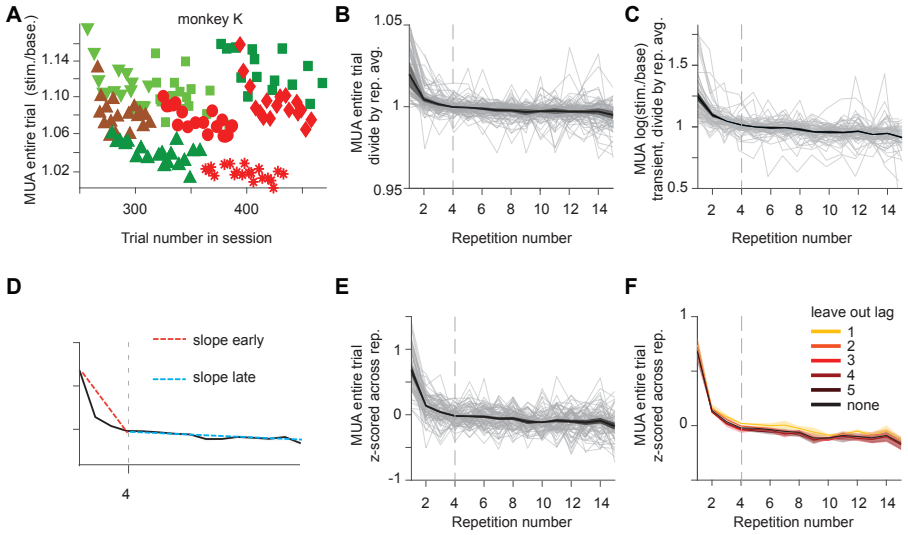

Figure 2.7 — Stimulus repetition effects on MUA responses show a characteristic shape. (A) Cross-site and trial time average responses relative to baseline to some example stimuli in one session of monkey K. Different color-symbol combinations indicate different stimuli. For some stimuli, a strong decrease from the first to the second presentation can be observed. (B) For each stimulus and animal, the change in response with repetition was computed by dividing each response by the average across repetitions. This was done for each site and session individually across the entire stimulation duration. Note that all sites were weighted equally, independently of their responsiveness to the stimulus, and that responses could be weak or even suppressed during larger parts of the trial (Figure 2.6). Cross-site, cross-session averages are shown for each stimulus (light gray lines), as well as the cross-stimulus, cross-animal mean and standard error across animals. Gray dashed line indicates the cutoff from the "early" to the "late" repetitions for slope estimations. Linear regression slopes were fit to the early (repetition 1:4) or late (repetition 5:15) responses (see Methods). Results are reported in the main text. (C) Same as (B), but for the initial response transient (0.05-0.15 s post-stimulus onset), log-transformed and averaged across sites and sessions before normalizing by the average response across repetitions. This de-emphasizes weakly responding sites compared to (B). (D) Sketch of the slope fitting procedure. (E) Same as (B) but for data z-scored across repetitions on the level of each individual site and session. (F) Analysis of the effect of lag. Same analysis as in (E) but with repetitions constituting a particular delay (1-5) from the previous stimulus presentation left out. Neither the early nor the late slope for the immediate repetition (lag 1) is significantly different from the slope where no lag has been removed (all $P \geq 0.15$).

only the stronger, transient response period was considered. Computed in this way, responses decreased by about 25% from the initial to later presentations on average. The effect size will be further discussed in the Discussion, section 2.5.2.

Both computations of effect size described so far consider only changes in mean, without considering the variance in response across repetitions. In Figure 2.7E, a final normalization, the z-score, is therefore considered. Z-scores are a commonly used measure of effect size and more robust than response ratios when responses are low. Here, for each stimulus, site and session individually, responses are normalized by subtracting the average response and then dividing by the standard deviation across repetitions. This analysis shows the degree of reliability of the repetition effect with respect to the within-session and within-stimulus single trial variance - even when averaging across the entire stimulation period and weighing all sites equally.

Fitting linear slopes to the early versus late repetitions (early till repetition 4, Figure 2.7D, see Methods) showed that responses decreased both for early and for later

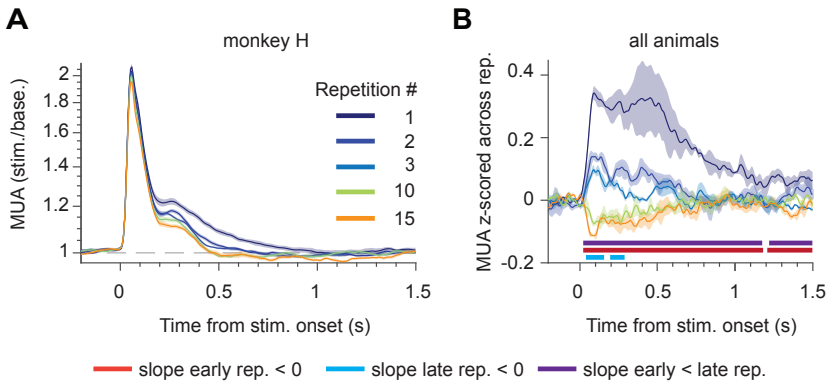

Figure 2.8 — (A) Cross-site, cross-session, cross-stimulus time-resolved MUA response for monkey H. Shaded region indicates standard error across sessions. (B) Cross-animal, cross-site z-scored time-resolved responses. Colored bars indicate multiple-comparison corrected tests for a negative early, late, or late vs. early slope across repetitions as defined in Figure 2.7D. Shaded area indicates standard error of the mean across animals. Horizontal bars indicate multiple-comparison corrected significance (Korn et al., 2004).

repetitions, but more strongly for early repetitions (all $P < 0.002$, permutation test, see Methods). Similar results were obtained using other normalizations. The differences between early and late slopes held consistently across stimuli (early slope z-scored data avg. -0.22, 98.7% of stimuli showed negative slopes across animals, late slope -0.01, 74.6% of stimuli) based on the slopes of the average repetition effect across sessions. Conclusions were also similar when based on the average of the slopes fit to individual sessions per stimulus. This indicates that the effect was consistent across stimuli and recording days. Note that a linear fit to early and late repetition trajectories was chosen because this is an easily implemented method that performs well on small amounts of potentially noisy data, such as single-trial data for individual stimuli and sessions. Furthermore, the current division can also be applied to gamma-band responses, which can show more complex repetition effects. It does not reflect an interpretation in the form of a genuine discontinuity between early and late repetitions. A log-log transformation did not fully linearize responses, indicating a very rapid decrease for initial repetitions (transformed data not shown).

Did the repetition effect derive only from immediate repetitions of the stimulus? Removing repetitions with a particular delay from the analysis did not appear to change the shape of the effect (Figure 2.7F). This held true even for removing immediate repetitions. Early and late slopes for the data excluding immediate repetitions were not significantly different from the slopes for all data (all $P \geq 0.15$). When considering averaged time-resolved responses, the repetition effect appeared rapidly after stimulus onset and lasted throughout the stimulation period (Figure 2.8A-B). The effect apparently decreased in strength around 750 ms. Around this time in the trial, MUA responses were frequently close to or even below baseline, potentially decreasing the sensitivity to detect repetition effects. Note that the short onset latency of the repetition effect excludes differences in eye movements or top-down cognitive modulations such as attention as the underlying cause for this effect (Lamme and Roelfsema, 2000; Poort et al., 2012).

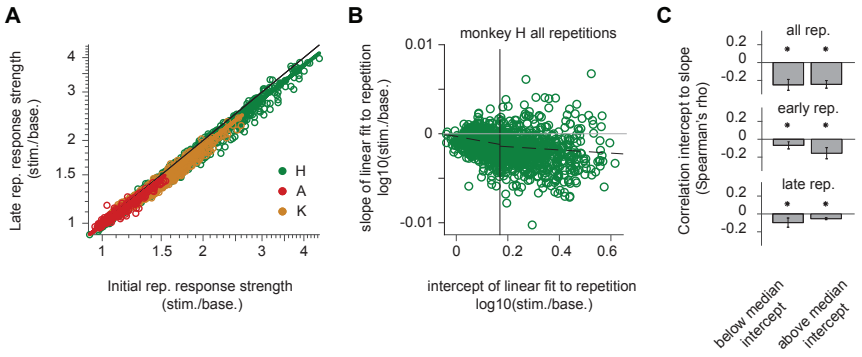

Figure 2.9 — Stimulus repetition effects and MUA response strength. (A) Initial response strength (stimulation/baseline, repetitions 1:4, see Figure 2.7) versus later response strength (repetitions 5:15) for individual site-stimulus combinations during the initial transient response period (0.05-0.15s post stimulus onset), averaged across recording sessions. Thick lines indicate robust regression linear fits for the individual animals, indicated by their respective colors. (B) Dependence of change in response (slope fit, see Methods and description in text) on initial response (intercept of linear fit) based on log(stim./base) responses, for an example animal, otherwise as (A). Vertical line indicates cutoff between the two data halves (median split), dashed black line indicates least squares fit to the data halves. (C) Average across-monkey correlation of change in response to initial response based on slope fits, with fits based on all, only the early or only the late repetitions. All correlations are tested for significance using permutation tests.

So far, the analyses focused on the average effects across sites. Next, we considered how the individual MUA site's response strength related to the repetition effect. Different models of stimulus repetition make different predictions about the dependence of repetition effects on the response strength. In particular, the output (and also the input) fatigue model would predict stronger reductions in responses with stronger drive. A rate-based sharpening model predicts stronger reductions with weaker responses. Figure 2.9A shows the average responses of each site to each stimulus for the early repetitions versus later repetitions, for the time period of the initial transient response (0.05-0.15 s after stimulus onset). Any repetition effects resulting from fatigue should be visible already during the initial response. For all animals, there was a clear positive relationship between initial and later repetitions (Spearman's rho 0.96-0.99 per animal), indicating that overall stimulus preference was maintained. Responses were also consistently stronger for the early presentations (Figure 2.9A). Furthermore, robust linear regression resulted in slopes below 1, which would indicate stronger decreases in response strength for stronger initial responses. However, note that this result may be driven by a floor effect, because stimulus-site combinations that show responses close to baseline to begin with cannot decrease their responses much. To test for a relationship between the response strength and its change with repetition, we used linear fits to all, or only the early or late repetitions (Figure 2.9B). The slopes of the linear fits were used as an indicator of change in response strength, and the intercept was used as an estimate of drive. Slopes and intercepts were estimated independently and in a cross-validated manner. Compared to a correlation between response magnitude and change in response, this analysis avoids circularity (see Methods for details). Correlations between slope and intercept estimates were computed for the lower and upper halves of the intercepts (i.e. response strength) independently (median split). This was done separately for each animal, and then averaged across

animals. The reasoning was that any correlation that persists in the upper half of response strength is unlikely to result from a floor effect. Figure 2.9B illustrates the correlation of upper versus lower halves of intercepts with the slopes, estimated based on all repetitions, for the animal with the largest range in response strength. The correlation was still negative for high response strengths. Across animals, correlations were consistently significantly negative. This held true when estimating slopes and intercepts based on all, or only the late or early repetitions (Figure 2.9C). Notably, the correlation between two random Gaussian variables is substantially higher than the correlation between a variable and the higher or lower median split of the other. Therefore, the correlation values here might underestimate the true underlying values. Though this property of the procedure is undesirable, we consider a too conservative estimate of the effect preferable to a "false alarm" due to the influence of a floor effect. The negative correlation between response strength and the repetition effect is in line with fatigue, but less so with sparsening proposals (see also Discussion).

2.4.4 Stimulus repetition effects on gamma responses differ from MUA in shape

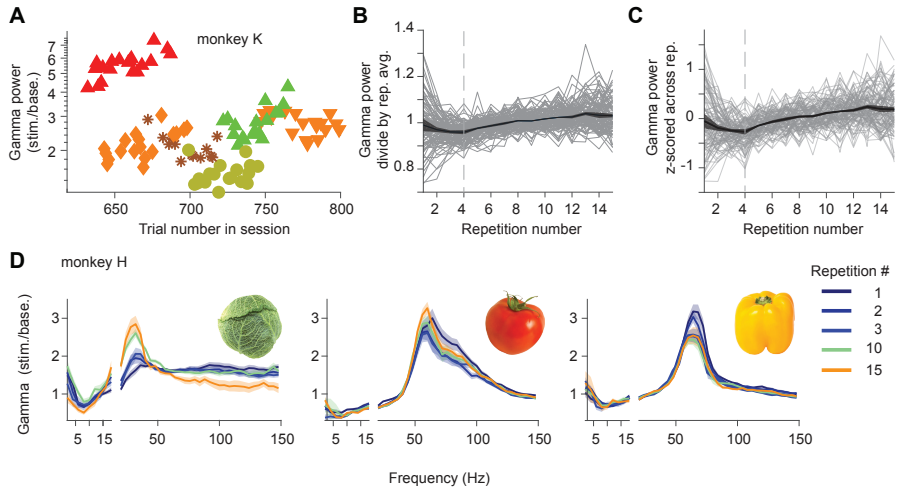

Figure 2.10 — Stimulus repetition effects on gamma-band responses. (A) Cross-site average gamma-band responses to some example stimuli in one session of monkey K. Gamma power is averaged around the largest peak defined per stimulus. Different color-symbol combinations indicate different stimuli. For some stimuli, a strong decrease from the initial to the later presentations can be observed. Some stimuli show an increase, across all or later repetitions. Note the log-scaling of the gamma power axis de-emphasizes changes for more strongly gamma-inducing stimuli. (B) For each stimulus and animal, the change in response with repetition for each site and session as in Figure 2.7B. Cross-site, cross-session changes with repetition are shown for each stimulus (light gray lines), as well as the cross-stimulus, cross-animal mean and standard error across animals. Gray dashed line indicates the cutoff from the "early" to the "late" repetitions for slope estimations. (C) Same as (B) but for data z-scored across repetitions on the level of each individual site and session. (D) Spectrally resolved cross-site, cross-session average repetition effect for example stimuli from monkey H (non-normalized). Shaded error represents the standard error over sessions. Frequency cutoffs and tapering as in Figure 2.5. Left panel shows increase, right decrease, middle combined pattern with repetition. Left panel also shows some broadband, high-frequency changes that might be related to changes in spiking activity or high-gamma.

We observed that repetition effects on gamma-band responses were more varied than effects on MUA. In particular, some stimuli showed decreasing responses with the first repetitions on average, whereas others did not. The direction of the repetition effect could change between the first few and later repetitions, which were dominated by response increases. Figure 2.10A illustrates a single-session example where different response patterns such as continuous increases, initial decreases and mixed responses can be observed. Figure 2.10B-C illustrates the average repetition effect for all stimuli and animals overlaid, based on, respectively, a division by the average response across repetitions per stimulus and site, or based on a z-score, similar to Figure 2.7. In Figure 2.10D, some spectrally resolved examples of the average repetition effect (non-normalized) across recording days for monkey H are shown.

We then quantified the repetition behavior of the gamma responses for early versus late repetitions. In Figure 2.11A, changes for early versus late repetitions are illustrated in a coordinate system. Each symbol corresponds to the average slope[3] of early versus late repetitions for a stimulus (across sessions and sites), for the z-scored data in Figure 2.10C. For example, a dot in the first quadrant corresponds to a stimulus that showed gamma increases for both early and late repetitions. A random distribution of slopes would cluster around zero and be equally distributed across all four quadrants. In contrast, the distribution of the data is biased towards the upper quadrants (late slopes tend to be positive) and somewhat less to the third quadrant (both early and late repetitions are related to decreases on average). A chi square test based on the pooled data across animals rejected the null hypothesis that the slope combinations are equally likely (chi squared value 39, df = 3, P = 1.7e-8). Furthermore, stimuli that showed a larger slope for initial presentations also tended to show a larger slope for later presentations (correlation between slopes r = 0.56, $P < 0.002$ based on a permutation test across animals).

For comparison, effects in MUA responses were evaluated in a similar manner, first for the entire trial period and then for the time period that the gamma-band power estimation is based on (0.5-1.5 s post-stimulus onset, Figure 2.11B-C). Results are in line with findings in the previous section: In contrast to the gamma-band slopes with repetition, MUA slopes are clustered in quadrants 2-3, more strongly than expected by chance (in both cases chi square ≥ 62, df = 3, P<= 2e-13). This indicates that especially early, and to a lesser degree late slopes, are typically negative. Regardless of the two time windows tested, there were no significant correlations between early and late slope for MUA responses (all Pearson's $r < 0.17$, all $P > 0.08$). Note that the pattern of MUA response slopes becomes even more concentrated when analysing the initial transient, where responses are strongest (84% ± 4% SEM of data in the third quadrant).

Comparing the pattern of gamma-band repetition effects with MUA repetition effects (computed on the same time window) showed that there were more gamma-band repetition effects in the first quadrant (early and late repetition effects are both increases) and more MUA repetition effects in the third quadrant (both early and late repetition effects are decreases, first quadrant difference -0.28, third quadrant 0.29, both $P < 0.008$ (Bonferroni corrected, two-sided permutation test), all other comparisons $P > 0.32$). The pattern of repetition effects is therefore both different from chance, and different between gamma-band and MUA responses. In particular, gamma-band responses are more likely to increase with repetition. Interestingly, stimuli that exhibited positive repetition effects already during early repetitions showed more positive rep-

[3] rather than the slope of the average response

etition effects also during later repetitions. Taken together, this indicates that gamma repetition effects may reflect two underlying and opposing processes. One process, which decreases gamma-band responses, may settle faster, such that later repetitions more strongly reflect the process that leads to gamma-band increases. This notion will be discussed further (Discussion sections 2.5.3 and 2.5.4).

Is there a relationship between the early (or late) slopes in MUA responses across stimuli with the early (or late) slopes in gamma-band power based on these average stimulus slopes? We tested these correlations for the gamma-band slopes with both MUA responses across the entire trial time and the same time window as used for gamma-band power estimation. We did not find significant correlations in either case, nor with slope estimates based on the initial MUA transient (all $r < 0.14$, all $P > 0.08$, uncorrected for multiple comparisons).

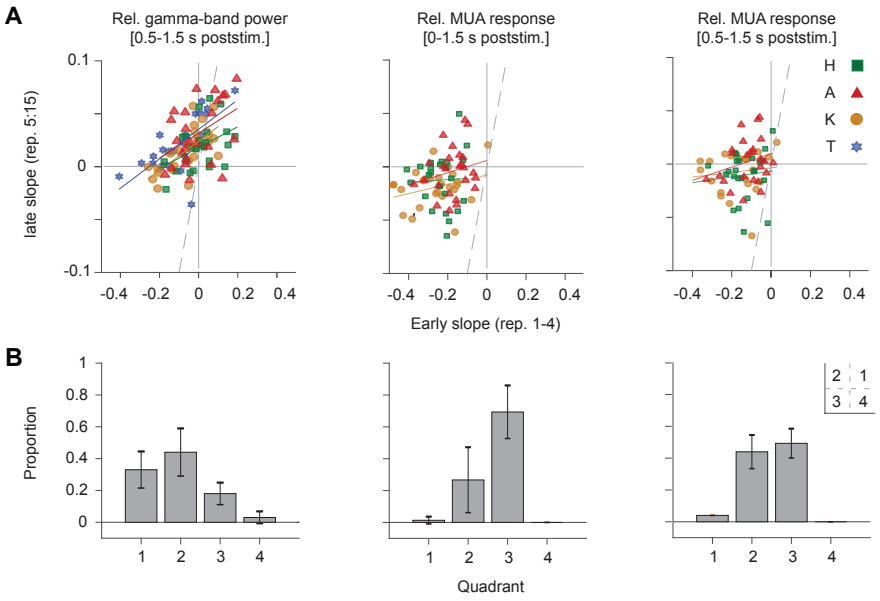

Figure 2.11 — Comparison of slopes fit to early vs late stimulus repetitions for gamma-band and MUA responses. (A) Each color-symbol combination represents the average early and late repetition slopes of a stimulus for z-scored data (see also Figure 2.10C). Different columns show results for gamma-band power, the MUA response (stim./base.) for the entire trial time as in Figure 2.7, and for the time used for the computation of gamma-band responses (0.5-1.5 s post stimulus onset). Colored lines indicate least squares linear fit of early to late slopes. Dashed line indicates the unity line. Note that the x- and y-axes have a different range to aid visibility of the late slope effects. (B) Proportion of stimuli occupying each quadrant. See inset for quadrant identity. Error bars denote standard deviation across animals. Statistical tests for both (A) and (B) are reported in the main text.

In Figures 2.5 and 2.6, it was shown that gamma-band responses to the natural images were stimulus specific and varied strongly in magnitude. Is there a relationship between the strength of the repetition effect in gamma-band responses to the overall response strength? Similarly to the analysis for MUA responses in Figure 2.9, Figure 2.12A shows the relationship between initial response strength (early repetitions) and later response strength. Overall correlations between early and late response strength were high, indicating that similar to MUA responses, overall stimulus preference was maintained (Spearman's rho 0.96-0.98). Since there was a mixture

of stimulus decreases or increases with repetition, responses clustered on both sides of the equality line. To test for a relationship between the response strength and its change, we used linear fits to all, or only the early or late repetitions, as in Figure 2.9 (see Methods).

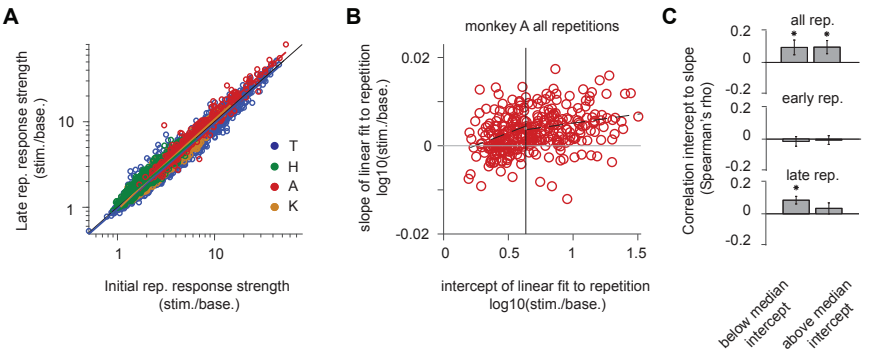

Figure 2.12 — Stimulus repetition effects and gamma-band response strength. (A) Initial response strength (stimulation/baseline, repetitions 1:4, see Figure 2.7) versus later response strength (repetitions 5:15) for individual site-stimulus combinations, averaged across recording sessions. (B) Dependence of change in response (slope fit, see Methods and description in text) on initial response (intercept of linear fit) based on log(stim./base.) responses, for an example animal, otherwise as (A). Vertical line indicates cutoff between the two data halves (median split), dashed black line indicates least squares fit to the data halves. (C) Average across-monkey correlation of change in response to initial response based on slope fits, with fits based on all, only the early or only the late repetitions. All correlations are tested for significance using permutation tests.

Based on fits to all repetitions, there was a small but significant positive correlation between initial gamma-band response strength (intercept) and its change with repetition (slope) both for the more weakly driven and the more strongly driven data half (Figure 2.12B-C). No such relationship was visible when considering only early repetitions, or for the more strongly driven stimulus-site combinations when considering only the late repetitions. As discussed in the corresponding section on the dependence of MUA repetition effects on response strength, the reported correlation values are likely underestimating the effect size due to the median split. Furthermore, note that data were log-transformed to reduce the influence of the strongest gamma-band responses on the correlation value (similarly to the MUA responses). Correlations were substantially higher without this log-transform. Overall, the pattern of responses indicates some positive correlation between stimulus-induced gamma-band response strength and a later increase. Fits for early repetitions are necessarily based on less data, which may have decreased sensitivity to detect an effect. An alternative explanation would be that gamma-band responses during early repetitions reflect a mixture of several processes. As shown in Figure 2.11, stimulus repetition led to mixed effects for early repetitions, and more consistently increases for later repetitions. If the mixed effects for early repetitions reflect the superposition of two different processes, one of which decreases gamma-band responses and settles faster, and another which increases gamma-band responses and is still ongoing during later repetitions, any dependence of the repetition effect on response strength could be masked for early repetitions (see also Discussion section 2.5.3).

2.4.5 Stimulus-specific changes in average gamma and MUA responses

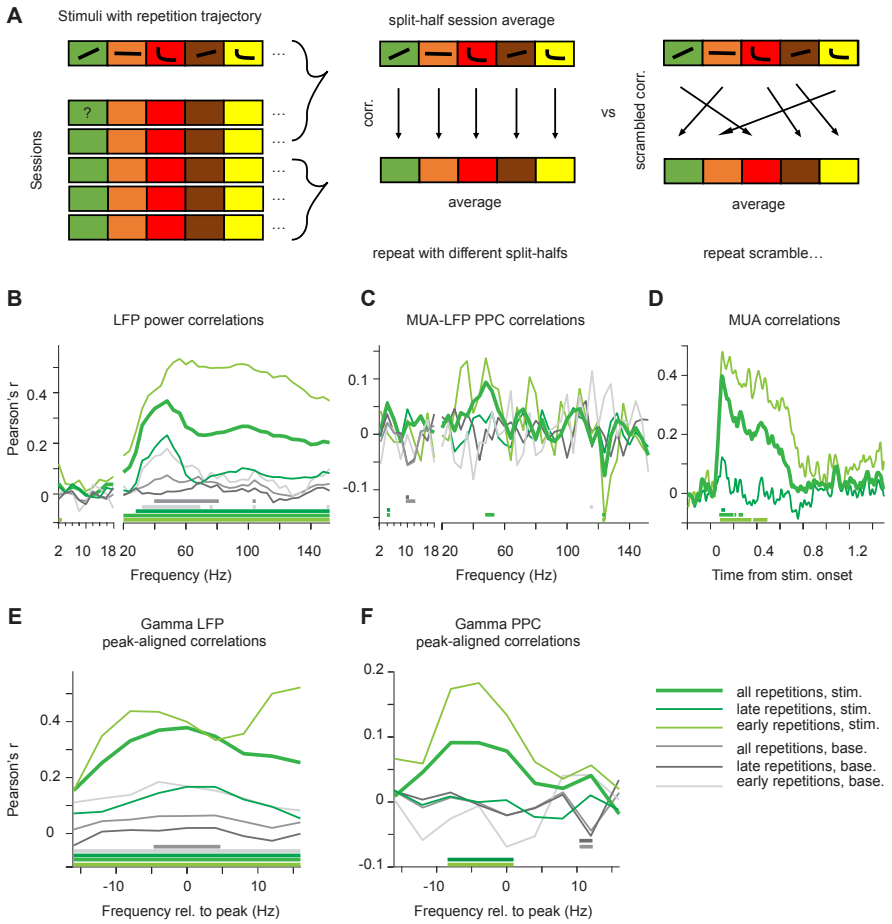

Figure 2.13 — Cross-session correlation of repetition effects indicate stimulus specificity. (A) Illustration of procedure (see Methods for full details). For each stimulus in each session, activity was z-scored across repetition for each recording site. Z-scored values were then averaged across sites, and across random halves of the recording sessions, yielding average repetition trajectories. Two split-half average repetition trajectories were correlated repeatedly across random partitions of sessions, and the correlation values compared to a permutation distribution based on randomly shuffling the stimulus identities. This was done for each frequency or time bin, and tests were multiple comparison corrected. (B) Spectra of LFP power correlations. (C) MUA-LFP PPC correlations. (D) Time-resolved MUA correlations. (E-F) Gamma-peak aligned versions of (B-C) respectively. (B-F): Thick green line indicates correlation for the entire repetition trajectory, based on responses relative to baseline (B,D,E) or on the stimulation period (C,F). Thin, bright green line indicates correlation values based on the first four stimulus repetitions. Thin, dark green line indicates correlation values based on the late repetitions (5-15). (B,C,E,F) Corresponding gray lines of different brightness indicate correlations during the baseline period. Horizontal bars at the bottom of each panel indicate significance, using multiple comparison correction (Korn et al., 2004).

Cross-session correlations indicate stimulus specificity. A stimulus-specific repetition effect should be reflected in correlations of the trajectory or shape of the repetition effect over days (Figure 2.13A). For example, a stimulus showing an increase on day

one should correlate more strongly with its own trajectory on the other days, than with trajectories of other stimuli (see Methods and Figure legend). We created z-scored repetition trajectories for each stimulus and each day, and ordered stimuli in a fixed order irrespective of appearance during the session, creating one vector per session. Session vectors were randomly split into two halves, and the averages of the two halves were correlated. Using a permutation test, we observed stimulus specificity in the form of between-session correlations of vectors of repetition trajectories that exceeded chance levels. For LFP power, these were strongest in the gamma-band and during visual stimulation (Figure 2.13B). Frequencies below 20 Hz showed no significant correlations with stimulus repetition. For correlations across all or just the late repetitions, a peak in the gamma-band was visible, but correlations remained significant also above 90 Hz. For correlations for early repetitions, the correlation was broader and more pronounced (differences between early only, late only or all repetitions not tested). This indicates that the significant correlation above 90 Hz could be partially the result of spiking (or high-gamma)[4] activity, which was stronger during initial trials. Furthermore, with the exception of late repetitions only, there were smaller yet significant correlations in the baseline activity. Interestingly, these were restricted to the gamma band.

Analyses that were aligned on the individual gamma-band peak per stimulus showed similar results (Figure 2.13E). For correlations including all or just early repetitions, the correlations seemed to be somewhat stronger for frequencies just below the peak. Possible reasons for this include asymmetries in the gamma-band spectra around the peak, or shifts in peak frequency with repetition.

Furthermore, the gamma-band peak-aligned correlations were of similar magnitude compared to the unaligned responses. A possible reason for this is that different stimuli show different correlation strengths over sessions. An unaligned spectrum will be biased towards the gamma-band peaks of stimuli showing the strongest correlations. Indeed, computing single-stimulus correlations on the average gamma-band peak aligned activity for all repetitions showed that correlations were positive on average (mean r = 0.29, SEM across animals 0.015), but could be close to zero - or even negative for some stimuli (12% of stimuli). In contrast, baseline activity was centered close to zero (mean r = 0.017, SEM 0.011, 45% of stimuli negative). Note that the approach of constructing vectors from all stimuli avoids the multiple comparison problem inherent in testing the many stimuli individually and allows to very specifically assess the stimulus-specificity of the correlation at the same time. Given that some stimuli contributed more to the overall positive correlation between session, we tested how many stimuli could be removed before correlations were no longer significant, for the average gamma-band peak aligned activity for all repetitions. To this end, we removed cumulatively stimuli from the overall repetition vectors, starting with the stimuli adding the most to the average correlation. Using multiple comparison correction, we found that up to the ten most positive stimuli (out of 25) could be dropped on average across animals before effects were no longer significant. In contrast, dropping even the single most positively correlated stimulus in the baseline reduced the correlation below the significance threshold. Is the vector-based correlation analysis driven by the stimuli with the strongest responses? There was a significant correlation between the stimuli with the most positive correlations and average gamma-band

[4]The "sculpting" of correlations from broadband more to the gamma-band for later repetitions was visible also in monkey T, whose ECoG implant did not show stimulus-evoked responses in the MUA. Also, note that with a cutoff of 4 trials, correlations were still significant for higher frequencies, with a cutoff of 5 trials, correlations were restricted to the gamma-band.

responses to the stimulus on the average across monkeys (r = 0.13, P < 0.05). Note, however, that this effect was driven by a single animal with a strong relationship between gamma-band response strength and correlation (monkey A r = 0.7, all others r = -0.1 to 0.14). We conclude that the cross-session correlations in repetition trajectories during stimulation are not the result of a few isolated stimuli showing correlations.

Next, we considered correlations for MUA-LFP phase locking (specifically, PPC, see Methods). MUA-LFP PPC correlations were lower overall compared to LFP power correlations, likely because single-trial estimates of PPC are noisy (Figure 2.13C and F). In the higher frequencies, there was a significant peak in the gamma-band when including all repetitions, and during the stimulation period only. In the lower frequencies, activity at 4 Hz during stimulation and at 10-12 Hz during the baseline showed significant positive and negative correlations. Note that negative correlations would mean that trajectories are anticorrelated between sessions more than expected by chance. The 4 Hz peaks are of similar absolute magnitude as these negative peaks. Whereas the peak in the gamma-band coincides with the peak in the LFP power correlations, there is no such correspondence for these lower-frequency findings. Given the small size of the PPC correlation values, these effects will not be considered further. The peak in the gamma-band was further supported by the gamma-peak aligned analyses, showing significant effects for all and also just the early repetitions. Interestingly, the correlations were shifted somewhat away from the PPC gamma peaks, similarly as in the LFP gamma-band aligned analyses. Note that the MUA-LFP PPC may suffer from several biases, which limit the interpretation on the PPC-based correlations. Firstly, PPC based on thresholded rates does not suffer from a spike-rate bias (Vinck et al., 2010b). However, the MUA signal employed here likely suffers from this effect, such that PPC values may be stronger for stronger rates, similarly to other measures of spike-field locking (Zeitler et al., 2006). Additionally, since the reliability of phase estimates will depend on the amplitude of the underlying activity, PPC estimates could also be affected by changes in power. Since both LFP power and MUA responses change, the effects on PPC and subsequently PPC repetition trajectories are difficult to estimate.

Finally, MUA activity showed significant correlations in the repetition trajectories. The time-courses of the effects are in agreement with Figure 2.8, but were only significant during the initial transient response after multiple-comparison correction. This indicates that inspite of a similar shape of the repetition effect between stimuli, correlations were significantly stronger between the repetition trajectories of the same stimulus than between stimuli. Overall correlation strength dropped when including all rather than just the initial 4 repetitions. Given the much smaller slopes of the repetition effect for later repetitions (Figure 2.7), effects for later repetitions were likely noisier, reducing overall correlation values. Single-stimulus correlations across the entire time window were r = 0.44 (SEM across animals 0.04) during stimulation and r = 0.06 (SEM 0.01) during the baseline period. There was no significant correlation between the correlation values and average MUA responses to the stimulus (r = 0.13, P = 0.14). Correlations of baseline-subtracted pupil responses averaged over the stimulus duration, computed based on all repetitions, showed no stimulus specificity (P = 0.49).

Multiple regression analysis of repetition effects. The correlation analyses above demonstrate that both MUA and gamma-band responses showed stimulus specificity in their repetition effects. The analysis was optimized to remove variance unrelated to stimulus repetition (by z-scoring across repetitions in a repetition sequence per stimulus, session and site, and by then averaging across session halves). Averaging

effects across session halves can reduce the effects of various sources of noise, including measurement noise and, given the design, the effects of varying the neighboring stimuli to a given stimulus, and effects of varying lag between repetitions. A different approach is to develop multiple regression models to fit single-trial responses directly. An advantage of this approach is that the effects of other variables of interest, such as reaction times, stimuli or pupil responses, can be incorporated. Any effect of stimulus repetition that remains will be independently predictive of the neuronal responses, at least to the degree that the contributions of other predictors are appropriately captured by a linear model. We therefore fit multiple linear regression models to single-trial MUA responses and single-trial gamma-band peak aligned responses, averaged across sites.

For the MUA responses, analyses focused on the initial transient response period (0.05-0.15 s post stimulus onset), though qualitatively similar results were obtained for analyses averaging across the entire stimulus time period, or across the time period for which gamma-band responses were computed (0.5-1.5 s post stimulus onset). Results for a fixed effect (monkey and session) model of transient, log-transformed relative MUA responses based on all 15 repetitions is shown in Table 2.3, including the proportion of variance explained and p-values based on an ANOVA (see Methods).

	Sum Sq.	df	Mean Sq.	F	Var. expl.	p
Stim. ID	10.8829	24	0.4535	2372.51	0.6156	0.0000
Lag	0.0783	1	0.0783	409.92	0.0044	0.0000
RT	0.0016	1	0.0016	8.39	0.0001	0.0038
Session:monkey	0.9777	46	0.0213	111.20	0.0553	0.0000
log(iRep)	0.2600	1	0.2600	1360.53	0.0147	0.0000
log(TrialInSess)	0.0201	1	0.0201	104.99	0.0011	0.0000
Stim. ID:Lag	0.0171	24	0.0007	3.72	0.0010	0.0000
Stim. ID:RT	0.0047	24	0.0002	1.03	0.0003	0.4161
Stim. ID:Monkey	3.2127	48	0.0669	350.19	0.1817	0.0000
Stim. ID:log(iRep)	0.0279	24	0.0012	6.09	0.0016	0.0000
Stim. ID:log(TrialInSess)	0.0384	24	0.0016	8.38	0.0022	0.0000
Error	2.1569	11285	0.0002		0.1220	

Table 2.3 — Results of regression model for MUA responses for the transient time window, log10(stim./base.) average across sites. Stim. ID = stimulus identity, RT = reaction time, iRep = repetition number, TrialInSess = total trial number in the session, Session:monkey = session number per animal.

Single-trial log-transformed MUA responses ranged between -0.005 and 0.25 (i.e. 0.99 - 1.78 fold increase with respect to baseline on this cross-site average, log-transformed before averaging). Stimulus identity (and its interaction with monkey, i.e. stimulus identity fit per animal) explained about 80% of variance. There was a significant effect of the log-transformed stimulus repetition number, capturing ca. 1.5% of variance, which showed significant modulation by stimulus identity (significant interaction term). The beta coefficients across stimuli were negative and are illustrated in Figure 2.14. There was a smaller and consistently positive effect of lag, indicating that more intervening stimuli led to somewhat stronger MUA responses (Figure 2.14). In addition, there was a significant effect of total trial number of the session, indicating a tendency for responses to decrease over the course of the session

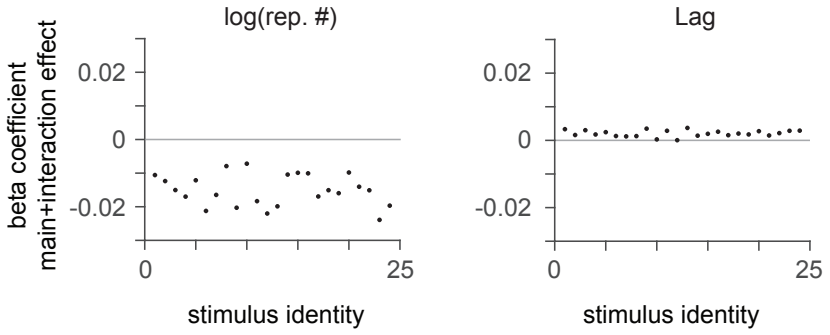

Figure 2.14 — Beta coefficients of the multiple linear regression model for MUA for the predictors log(repetition number) and lag, resolved by stimulus identity. Stimulus repetition had negative, and increasing lag smaller, positive effects on MUA responses. For comparison, beta values of stimulus identity ranged between -0.08 and 0.08. This model did not include a higher order monkey:stimulus:repetition interaction, but note that models based on individual animals were qualitatively similar.

(mean beta = -0.003, range -0.015 to 0.006 across stimuli). MUA responses tended to be higher when subsequent reaction times were faster (RT in seconds, beta = -0.006, no significant interaction with stimulus identity). Additionally, there was a significant effect of session, which as a categorical predictor captured individual offsets between recording sessions. The present model based on log-transformed repetition numbers and all repetitions was chosen to capture the repetition effect in MUA responses in a single model, weighing the effect approximately by the observed shape. The qualitative pattern of results, and the significant effect of repetition number, remained when analyses were performed for individual animals, for the non-log-transformed repetition number, when including the pupil response for animals where it was available, and when models were fit only to the initial 4 or only on the later repetitions (log-transformed or not). Models fit on only the later repetitions still showed significant effects of repetition. These explained less variance by a factor of ca. 10 compared to models based on early or all repetitions.

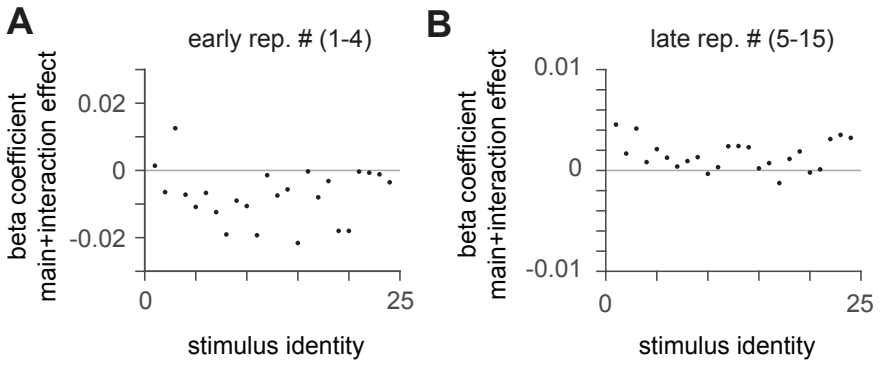

Figure 2.15 — Beta coefficients of the two multiple linear regression models (early or late repetitions only) for the predictor repetition number, for the gamma-band activity. The stimulus repetition predictor changed qualitatively between early and late repetition models, from largely negative to more positive, but smaller effects.

	Sum Sq.	df	Mean Sq.	F	Var. expl.	p
Stim. ID	78.1083	24	3.2545	620.78	0.5170	0.0000
Lag	0.2156	1	0.2156	41.13	0.0014	0.0000
RT	0.0083	1	0.0083	1.58	0.0001	0.2090
Session:monkey	11.6089	65	0.1786	34.07	0.0768	0.0000
iRep	0.2756	1	0.2756	52.58	0.0018	0.0000
log(TrialInSess)	1.0562	1	1.0562	201.46	0.0070	0.0000
Stim. ID:Lag	0.1803	24	0.0075	1.43	0.0012	0.0789
Stim. ID:RT	0.1255	24	0.0052	1.00	0.0008	0.4649
Stim. ID:Monkey	34.6634	72	0.4814	91.83	0.2294	0.0000
Stim. ID:iRep	0.3187	24	0.0133	2.53	0.0021	0.0001
Stim. ID:log(TrialInSess)	0.7098	24	0.0296	5.64	0.0047	0.0000
Error	23.8225	4544	0.0052		0.1577	

Table 2.4 — Results of regression model for peak-aligned gamma-band responses for the first 4 repetitions, log(stim./base.) average across sites. Stim. ID = stimulus identity, RT = reaction time, iRep = repetition number, TrialInSess = total trial number in the session, Session:monkey = session number per animal.

	Sum Sq.	df	Mean Sq.	F	Var. expl.	p
Stim. ID	212.3025	24	8.8459	1895.90	0.5238	0.0000
Lag	1.4217	1	1.4217	304.71	0.0035	0.0000
RT	0.0047	1	0.0047	1.00	0.0000	0.3167
Session:monkey	35.8398	64	0.5600	120.02	0.0884	0.0000
iRep	0.2251	1	0.2251	48.25	0.0006	0.0000
log(TrialInSess)	1.6012	1	1.6012	343.17	0.0040	0.0000
Stim. ID:Lag	0.3368	24	0.0140	3.01	0.0008	0.0000
Stim. ID:RT	0.1049	24	0.0044	0.94	0.0003	0.5500
Stim. ID:Monkey	98.4373	72	1.3672	293.02	0.2428	0.0000
Stim. ID:iRep	0.2026	24	0.0084	1.81	0.0005	0.0090
Stim. ID:log(TrialInSess)	1.1674	24	0.0486	10.43	0.0029	0.0000
Error	53.7038	11510	0.0047		0.1325	

Table 2.5 — Results of regression model for peak-aligned gamma-band responses for the late (5-15) repetitions, log(stim./base.) average across sites. Stim. ID = stimulus identity, RT = reaction time, iRep = repetition number, TrialInSess = total trial number in the session, Session:monkey = session number per animal.

We then performed similar analyses for log-transformed gamma-band responses, with single-trial responses estimated around the gamma-band peak per stimulus, averaged across sites showing this peak in their spectrum (see Methods). Single-trial log-transformed gamma-band responses ranged between 0.015-1.36 (i.e. 1.03-22.9 fold increase with respect to baseline on this cross-site average, log-transformed before averaging). In contrast to the MUA responses, and in line with the previous analyses on the gamma-band, models fit to the early compared to later repetitions showed a qualitative change in the repetition effect, such that two models will be presented here. Both models were based on non-log-transformed repetition number. Both models showed small but significant effects of stimulus repetition. In the model for early rep-

etitions (Table 2.4 and Figure 2.15A), repetition effects were predominantly negative. In contrast, effects were smaller and predominantly positive for late repetitions. As for the MUA response model, most variance was explained by stimulus identity (ca. 75% when taking differences between animals into account) and offsets between individual sessions. There was no significant effect of reaction times in either model. Increasing lag tended to have positive effects on gamma-band power overall (early rep. model: mean beta = 0.005, range -0.004-0.016; late rep. model: mean beta = 0.009, range 0.0015-0.020). For both models, there was a significant effect of the log of the total trial number in the session, which tended to be positive, and explained more variance than the stimulus repetition number (both models highly similarly: mean beta = 0.05, range -0.002-0.11)[5]. It is possible that this effect is explained by a general shift in arousal with the session. An interesting alternative is that repetition effects could to some degree "spill-over" between stimuli, possibly between more similar ones, even with substantial delay. This notion is supported by significant interactions between stimulus identity and the trial number in the session, which could indicate that some stimuli are influenced more by other stimuli than others. Both the qualitative pattern of results and the significant effect of repetition number, were unchanged when analyses were performed for individual animals, and when including the pupil response for animals where it was available.

2.5 Discussion

A set of 25 natural images was presented repeatedly in a pseudorandom fashion to monkeys performing a change detection task. The natural images induced a wide range of MUA and gamma-band responses in V1. Upon repeated presentation within a recording session, both MUA and gamma-band responses showed stimulus repetition effects.

There are several motivations behind measuring repetition effects in LFPs and spike-field synchronization and comparing them to effects in spiking activity. First, it has been proposed that synchronization of responses may implement adaptive effects of stimulus repetition, by increasing the effectiveness of interareal or within-area coordination (Brunet et al., 2014; Gotts et al., 2012; Gotts, 2003). Second, stimulus repetition will affect networks of neurons, and possibly properties of the network (e.g. balance between excitation and inhibition or interneuronal correlations, De Baene and Vogels 2010; Natan et al. 2017; Gutnisky and Dragoi 2008), which may be reflected in altered gamma-band responses. Third, gamma-band responses can be measured non-invasively in humans using MEG, and can be more strongly correlated with the fMRI signal than spiking activity, thereby providing a crucial link between human and animal studies of plasticity (Niessing et al., 2005; Logothetis and Wandell, 2004; Ekstrom, 2010; Maier et al., 2008; Thomsen et al., 2004; Viswanathan and Freeman, 2007; Nir et al., 2007; Scheeringa et al., 2016; Bartolo et al., 2011).

[5]For both MUA and gamma-band responses, a log(TrialInSess) predictor outperformed predictors of total trial time or linear predictors of trial number in session.

2.5.1 Stimulus specificity of repetition effects and generalization to natural images

A central question of the current experiment was whether the previously observed repetition effects with grating stimuli (Brunet et al., 2014) were stimulus specific. If a stimulus is presented repeatedly with different neighboring stimuli and in different parts of a session, is the repetition effect preserved? Furthermore, we asked whether the observed phenomena would generalize to a broader range of more naturalistic stimuli that were initially novel to the animals, rather than using stimuli that were highly overtrained and designed to generate strong responses. We indeed found that there was stimulus specificity of the repetition effect in the MUA responses, as well as gamma-band LFP and PPC responses, but not in responses in lower frequency bands (Figure 2.13). Cross-session correlation values were stronger for initial repetitions, which induced stronger changes overall, than for later repetitions, probably reflecting changes in SNR (see below for a discussion of the shape of repetition effects). For gamma-band power, smaller but significant correlation values were also observed for the baseline period. This could reflect some stimulus-specific memory trace, since stimulus repetitions had a constrained lag, such that baseline activity was associated with a specific stimulus well above chance level. However, sequentially dropping out stimuli that contributed to this correlation showed that baseline effects may have been limited to some very few stimuli, in contrast to effects during the stimulus period. For both MUA and gamma-band analyses, correlation values were below $r = 0.5$ (sometimes much lower, depending on the exact signal and range of repetitions in question). Similarly, regression analyses showed that stimulus-specific repetition effects explained significant, but small amounts of variance for both MUA responses and gamma-band power when compared to the size of the effects of stimulus identity. To better understand these effect sizes, one should consider several aspects of the current design, namely 1) the choice of stimuli, 2) the within-session and between-session design and 3) resulting challenges for the analysis, of synchronization in particular.

1. In order to test stimulus specificity and the generalization of the effect to more naturalistic conditions, we chose a set of 25 different natural images, with previously unknown response properties in V1, especially with regard to gamma-band responses. This allowed the present conclusion that repetition effects generalize to a broader range of stimuli. It also induced strong variance in the response strength, with stimuli inducing response increases from tens of percent to thousands of percent in the gamma-band and strong differences in MUA responses (Figure 2.6). Such variance is interesting as such and in this specific experiment, was a part of the question. However, it goes against the basic design principle in experiments to minimize variance of any effect that is not the main effect of interest. For example, Vogels (2016) reviews studies that typically first identify two stimuli driving a cell equally well, before proceding with repetition tests. In regression models, reducing the overall variance induced by different stimuli will increase the overall percentage of variance explainable by a repetition effect.
2. Both the within-session and across-session design introduced “noise” or variance to the stimulus repetition effects, by randomly changing both the immediate neighbors and timing between repetitions (lag). Stimulus repetition effects are sensitive to such contextual variation (e.g. Solomon and Kohn 2014 and

Kim et al. 2019), which will reduce the detectability of stimulus specificity in this design.

3. A further consequence of 2. is that responses could not always be averaged, and a lot of analyses required single-trial, and thereby noisy, estimates of responses. For example, MUA-LFP PPC, though unbiased by trial number, benefits strongly from increasing trial numbers in terms of accuracy (Vinck et al., 2012, 2010b). The effect size for MUA responses will be discussed further in the next section.

To summarise, both MUA and gamma-band responses showed stimulus-specific repetition effects for natural images, initially novel to the animals. The effects were detectable and therefore to some degree robust against variations in neighboring stimuli in time and the precise timing of the repetitions, both within a session (lag) and between sessions (position in the session).

2.5.2 Effect size, timing and shape of MUA repetition effects

Effect size of MUA repetition effects. MUA responses decreased most rapidly for initial presentations, and more weakly for later presentations. On average across the entire stimulus duration and all recording sites regardless of their stimulus responsiveness, repetition decreased responses only by a few percentage points (Figure 2.5B). However, overall MUA responses to the stimuli were not increasing much over the gray background baseline when computed this way. There are likely multiple reasons for this low responsiveness, four of which will be discussed here. 1) Responses were low, or even dropped below baseline, in later parts of the trial. If responses were only dropping toward baseline, the underlying reason for this may be the relatively long trial durations and subsequent adaptation. The drop of activity below baseline, which could occur relatively rapidly (see Figure 2.5), indicates that inhibitory mechanisms may also be at play. This is not unlikely because 2), due to the size of the stimuli, individual V1 sites likely exhibited surround suppression. Surround suppression occurs both in V1 and in its inputs, such that large stimuli may induce both weaker initial and weaker later MUA responses. The topic of surround suppression for large, colored stimuli will be investigated in Chapter 4 and then discussed at length. 3) Stimuli were not designed to drive V1 or particular V1 sites strongly. In contrast, a long history of visual neuroscience has relied on identifying stimuli that drive individual cells well (Solomon and Kohn 2014, Wissig and Kohn 2012, see General Discussion). 4) The recording arrays yielded MUA but not single-unit responses. Since different neurons may show different repetition effects, it is conceivable that these average out to some degree. In addition, given that arrays are in long-term use, the overall SNR in chronic recordings tends to be lower than in acute recordings.

Notably, analyses focused on the initial transient response period and biased toward more strongly responsive sites showed larger effect sizes (ca. 25 percentage points on average, Figure 2.5B). For comparison, in V4, a single, rapid stimulus repetition resulted in a 30 percent reduction in response rate for optimally driven cells (Wang et al., 2011). We therefore conclude that V1 may show repetition effects at least on a similar order of magnitude compared to mid-level and higher areas. Given differences in design to previous studies, a more accurate comparison will require multi-area recordings under matched conditions (e.g. stimulus eccentricity, duration, size with respect to the receptive field).

Timing of MUA repetition effects. Time-resolved analyses showed that the repetition effects had a rapid onset (Figure 2.8). This excludes factors such as top-down modulation by attention or eye movements as the main contributing factors to the repetition effect (Lamme and Roelfsema, 2000; Li et al., 2004; Lee and Nguyen, 2001; Supèr et al., 2003; Poort et al., 2012). In terms of absolute or percentage changes, the effects were strongest sometime after the initial transient (> 100 ms). This timescale is reminiscent of the timing of the largest effects in IT for immediate repetitions, and the timing of response reductions through long-term exposure in V2 and IT (Vogels, 2016; Huang et al., 2018; Freedman et al., 2006). Such timing has been interpreted as an indication that the repetition effects depend on changes in recurrent processing within or across brain areas (Vogels, 2016; Huang et al., 2018). This interpretation is also supported by the present finding of changes in gamma-band responses, which were computed over a late part of the trial, and are thought to depend on recurrency (see General Discussion 6.4). We also observed a change in the very early response strength, that was smaller in percent change, as can also be observed in higher-order areas (Vogels, 2016; De Baene and Vogels, 2010). Interestingly, when assessed with a z-score, which scales the size of the effect with its reliability, early response changes had an effect size similar to the later changes. For a neuron in a downstream area receiving many inputs, the effect may therefore be of similar detectability early on. Since relative response changes are in frequent use compared to z-scores, a comparison of this perspective across areas is currently not possible. Such an immediate effect of stimulus repetition on MUA responses is in line with both an output fatigue and an input fatigue account of stimulus repetition.

*Duration of MUA repetition effects.*We also observed the repetition effect occured irrespectively of lag (Figure 2.5F, regression models indicated a small, modulatory effect), indicating that it was robust against several intervening stimuli and had a memory of several seconds[6]. This exceeds previously reported durations of history effects in rodent and cat V1 (Kim et al., 2019; Nikolić et al., 2009; Patterson et al., 2013). It has been suggested that the duration of repetition effects is related to their intensity (duration-scaling, Bao and Engel 2012). Since the current stimulus durations were on the order of several seconds, which is longer than in most other studies, it is possible that the history effect observed here resulted from such aspects of the design. Patterson et al. (2013) showed that for large stimuli, V1 peak response changes under anesthesia showed little dependence on stimulus durations of 0.4, 4 and 40 s. In contrast, time to recovery from adaptation was clearly dependent on stimulus duration. It is therefore possible that the history effects in the present data reflect the stimulus duration. However, Patterson et al. (2013) showed that effects of 4 s stimulation were recovered after 4 s of delay (without interfering stimuli). It is therefore also possible that the memory effects seen here rely on plasticity mechanisms that are only at work during the awake state. Recently, a similar stimulus duration to the current case (3 s) was tested and directly compared to shorter stimulus durations in IT of awake animals (Kuravi and Vogels, 2017). It was found that the number of repetitions, rather than stimulus duration, was determining repetition effects. However, only immediate repetition effects were explored, such that it is conceivable that repetition effects did not affect MUA responses more strongly, but did leave a longer memory trace nonetheless, in line with Patterson et al. (2013). The question of memory in V1 for stimulus repetition will be explored further in the next chapter.

[6]The term memory here refers to a persistent effect in the system, i.e. V1, and not a memory in the psychological sense.

Shape of MUA repetition effects. The decrease in MUA responses with repetition was stronger for initial repetitions compared to later repetitions. This is reminiscent of the shape of repetition effects for repetitions within a trial in IT that are stimulus-specific (Sawamura et al., 2006; Kaliukhovich and Vogels, 2014), and on a very different time scale, repetition effects in IT for cases with very many intervening trials (Li et al., 1993). In monkey V1, rapid decreases for the initial repetition *within* a trial have been shown, although the stimulus-specificity of the effect is unclear in this case (Westerberg et al., 2019). Given the design in this particular study, the first stimulus in a sequence will have experienced a larger inter-stimulus-interval to the preceding stimulation, and thereby be disadapted more strongly (or as a different perspective, more strongly adapted to the gray background). In contrast, in the current design, repetitions occured between trials, such that any general form of disadaptation should affect all trials equally. We will turn to a discussion of potential underlying reasons for the observed shape after also considering the repetition effects in the gamma-band.

2.5.3 Shape of gamma repetition effects suggests two underlying processes

Using the same fitting procedures as for MUA responses, we found that the average repetition effect for gamma-band responses was more diverse across stimuli than the average MUA effect. Whereas MUA responses were dominated by response decreases, gamma responses could show both increases or decreases during the initial repetitions. During later repetitions, most stimuli showed increases. On average across sites and sessions, these changes were on the order of a few (ca. 10) percentage points increase in response for the late repetitions, or about 0.5 z-scores. Effects could be more pronounced but were also more varied for initial repetitions. Interestingly, there was a positive correlation of r = 0.56 between the slope for initial and later repetitions on average.

The shape and distribution of the repetition effects in gamma-band power suggests that there may be two repetition-related processes that affect gamma-band power in opposing ways, with different time constants, and different stimulus-dependence. From the MUA responses, at least one repetition-related process with stronger effects for initial presentations can be inferred. It is possible that such a process also decreases gamma-band responses, for example by reducing overall drive in the network (see General Discussion 6.2). At the same time, the dominance of increases with late repetitions, in line with previous work (Brunet et al., 2014), suggests that stimulus repetition can increase gamma-band responses. This process remains strong for more repetitions, showing a log-linear relationship on the scale of hundreds of binned repetitions, with the first bin containing about 4 times the repetitions as tested here in total (Brunet et al., 2014). (This feature is used as part of the experimental design in the next chapter.) We can therefore speculate that there are two processes, one fast and one slow in terms of repetition numbers to asymptote. Depending on the relative strength of the first and second process, the repetition effects could add up to both response decreases and increases (or also no apparent change) in the first repetitions, whereas later repetitions may be dominated more strongly by the second process. This scenario is supported by the correlation between initial and later slopes of repetition effects in the gamma-band response, which could result from the second process dominating already initial responses when it is strong (see Figure 2.16 for illustration).

In contrast to the present study, Brunet et al. (2014) did not report an initial decrease with stimulus repetition, which in that study coincided with the beginning of

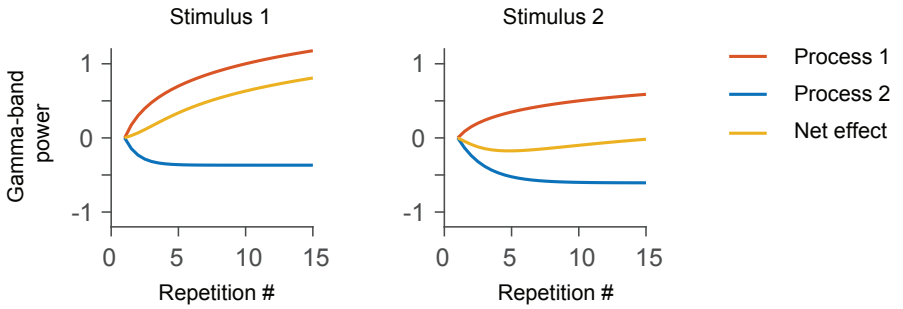

Figure 2.16 — Illustration of opposing effects of two underlying, stimulus-dependent repetition processes for two example stimuli. Depending on the relative strength of the two processes, the net repetition effect can be positive or negative for the initial repetitions. Gamma-band power values are arbitrary.

the recording session. Based on the current results, this could have several reasons. 1) The first few presentations in a recording session were "special" and masked the initial decrease, potentially if the animal made mistakes on the first few trials, or experienced similar stimulation before the start of the recordings. 2) The stimulus employed, a moving, colored grating, was optimized to induce strong gamma-band oscillations. Therefore, the second, slower process that increases gamma-band responses could dominate already during the first repetitions, or balance to a net effect close to no change. Notably, Figure 1J of Brunet et al. (2014) shows little response change in the first few repetitions.

2.5.4 Repetition effects, fast and slow

The shape of the repetition effect in the MUA is decreasing monotonically, such that a single underlying mechanism may be the most parsimonious explanation. There is also some psychological appeal to the notion that a single experimental manipulation results in a single effect. However, adaptation or stimulus repetition paradigms can induce plasticity through a variety of mechanisms, and on different time scales, which may be involved simultaneously (Solomon and Kohn, 2014; Patterson et al., 2013). Based on theoretical considerations of the "predictive coding" framework, it has been suggested that stimulus repetition effects involve (at least) two processes, one of them acting very rapidly (Auksztulewicz and Friston, 2016; Friston, 2005, 2008). The fast process was suggested to reflect a decrease in precision of prediction errors after the presentation of a novel stimulus (confidence about prediction is lost with a novel stimulus, and recovered with first repetition). The slower process is interpreted as a cumulative reduction in error signaling. In line with this notion, Garrido et al. (2009) showed that there is a faster and a slower process reflected in EEG responses to repeated tones.

The initial, rapid response decrease could also be seen as a response *increase* for novel stimuli. The distinction between a novelty increase and familiarity decrease requires specific designs and can show stimulus- and brain area (and possibly species) dependence (Amado and Kovács, 2016; Vinken et al., 2017; Kaliukhovich and Vogels, 2014). The present experiment was not designed to distinguish these possibilities. An

interpretation in terms of a novelty response[7] allows another functional interpretation of an initial fast process. The relatively indiscriminate response increase across recording sites observed here is in line with the idea that when a change occurs, a robust response for detection may be more important than accurate but possibly slow processing. In a natural environment, a sensory stimulus may then be sampled repeatedly for refined processing (e.g. with eye movements for vision, Zhaoping 2019, or breathing for olfaction, both processes occur rhythmically several times per second). From a mechanistic perspective, it is known that some forms of synaptic depression can act upon a single stimulation (Abbott, 1997), whereas other processes may accumulate over a number of repetitions (Fioravante and Regehr, 2011; Malenka and Bear, 2004). An initially strong response (and strong decrease with a first repetition) is therefore in line with change detection proposals, possibly based on input fatigue mechanisms. A second process could be engaged to refine responses to the repeated stimulation, for example through synchronization-based sharpening (Brunet et al., 2014) and/or improved discriminability of fine differences (Hansen and Dragoi, 2011; Wang et al., 2011).

From the perspective of gamma-band responses, the notion of (at least) two competing processes with different time scales (and with different strength depending on the stimulus, Figure 2.16) may reconcile some apparent contradictions in the literature. Some previous studies using *single* repetitions have reported decreases in the gamma-band with repetition (Friese et al., 2012; Kaliukhovich and Vogels, 2012). Since the achromatic stimuli used may not have induced strong gamma-band responses, the decreasing, fast process observed in the current paradigm may have dominated the results in Friese et al. (2012). In the case of IT, the activity resided in what could be considered high-gamma (Kaliukhovich and Vogels, 2012). High-gamma activity has recently been shown to be dependent on NMDA receptors, which are implicated in learning and plasticity, and changes in dendritic processing (Bartoli et al., 2019). Therefore, such a finding is in line with the notion of a rapid novelty-related response and subsequent synaptic changes. The more broad-band stimulus-specific correlations for initial repetitions (Figure 2.13) further support the notion that initial repetitions may induce high-gamma activity (see also next chapter). Hansen and Dragoi (2011) and Wang et al. (2011) reported gamma-band increases for a single repetition of a grating stimulus, when compared to a previous presentation of a random dot stimulus, and restricted to gratings that generated strong gamma-band responses to begin with. Given the design, i.e. using the same grating stimuli often, stimuli were only novel at the very beginning of a session. Moreover, the restriction in the stimulus set may have masked an initial decrease in gamma-band activity, with a similar reasoning as for Brunet et al. (2014). Such a design could therefore unmask the second process even with a single, within-trial repetition paradigm.

Interestingly, it is possible that both the rapid and the slow process are prevalent (in some related form) across species and systems. In the insect and rodent olfactory system, high-frequency oscillatory responses increase over the course of a few to tens of repetitions (van Wingerden et al., 2012, 2010; Laurent and Davidowitz, 1994; Laurent et al., 1999; Stopfer and Laurent, 1999). In insects, the firing rates upon stimulus repetition decrease most strongly for the first repetition, whereas the oscillatory responses show a more steady increases for about 10 odor repetitions. It is possible that differences between insects and mammals, including the lower frequency of the odor-induced oscillation, prevent an overlap of the two processes in the LFP responses. In

[7] This also applies to temporally local novelty, even if stimuli have been seen on previous days.

rodents, careful inspection of van Wingerden et al. (2010) shows that the initial 2-3 repetitions are characterized by a small decrease or flat response in the gamma-band on average (Figure 3A-B). Interestingly, a strong decrease in response for the first repetition of an odor stimulus has also been shown in zebrafish telencephalon, where the imaging technique used prevented the study of oscillations (Jacobson et al., 2018).

2.5.5 Stimulus repetition versus prolonged exposure

The preceding discussion assumed that stimulus repetition underlies the observed effects, rather than an increase in total exposure duration to the stimulus. In other words, do we need to repeat a stimulus to see effects, or would prolonged exposure suffice? It is likely that prolonged exposure results in a third form of plasticity, a form of adaptation that may begin at the level of the retina with genuine output fatigue (Kohn, 2007; Solomon and Kohn, 2014). Jia et al. (2011) showed that prolonged exposure to a grating stimulus decreased gamma-band responses specifically to this stimulus, and could increase responses to other stimuli. In line with this, in a following chapter, we will show that prolonged exposure reduces responses to colored stimuli (Chapter 5), which was also observed for some of the natural stimuli (data not shown). For grating stimuli, a within-trial increase in gamma-band responses has been observed that matched the hazard rate of trial duration (Lima et al., 2011). However, this was not observed here, possibly because the hazard rate was kept more linear, possibly because of fundamentally different ways these grating stimuli adapt compared to colored stimuli (Spyropoulos et al., 2019). Kuravi and Vogels (2017) showed that in IT, repetition effects strongly depended on an actual repetition, but showed no detectable difference between a preceding stimulus of 300 ms vs 3 s duration in the MUA response. A stimulus repetition (which entails a break in stimulation) may allow for disadaptation of upstream signals, including receptors. Repetitions also bear closer resemblance to sampling behavior during vision. It therefore seems that the stimulus repetition effect is indeed related to repeated presentation per se. Compared to prolonged stimulation, repetition also entails a relatively unpredictable, typically flashed stimulus onset that induces a transient response, which could be important for repetition effects. Although eye movements also cause strong shifts in visual inputs, it is unclear to what degree such a flashed onset sufficiently mimicks processes in natural vision.

2.5.6 Dependence of repetition effects on initial response strength

We observed a dependence of both MUA and gamma-band effects on the initial response strength, in opposing directions. In order to avoid circularity, response strength and change in response with repetition were estimated using the intercept and slope, respectively, of a linear, cross-validated fit to the repetition trajectory. In order to avoid conclusions based on a floor effect due to poorly responsive sites, we then tested correlations between response strength and change with repetition for two data halves independently, median split by response strength. This is a conservative procedure that can underestimate the true underlying correlation present, in cases where there is no floor effect.

Stronger MUA responses for a given stimulus and site tended to decrease more strongly with repetition. The dependence of MUA repetition effects on response strength held true both for early and for late repetitions. This finding, similar to the timing of the MUA repetition effects, is in line with both output fatigue and input fa-

tigue proposals. It is not in line with a rate-based sharpening account - unless some very highly responsive and selective cells carry the effect, which may not be visible in MUA responses. Similar results have been reported previously for single units in IT for paradigms relying on short-term rather than long-term plasticity (McMahon and Olson, 2007; De Baene and Vogels, 2010; Miller et al., 1993). De Baene and Vogels (2010) explicitly tested for a sharpening account in IT for both spiking and high-gamma activity, and found that there was scaling rather than sharpening of responses, independently of the task the animals were performing. Wissig and Kohn (2012) reported that for a prolonged stimulation paradigm in anesthetized V1, response reductions were stronger for more strongly driven cells. However, for large stimuli, response reductions were weaker overall than for small stimuli, such that the dependency was relatively small for large stimuli (similar to the correlation values observed in this experiment). Given that in the present study, 1) responses were low overall, 2) the role of lag between repetitions was small, and 3) the effect was stimulus-specific, the overall evidence tends to support input fatigue (e.g. synaptic plasticity) rather than output fatigue. This would imply that stimulus repetition involves a network effect, which is also supported by the change in gamma-band responses with repetition.

Stronger gamma-band responses for a given stimulus and site tended to increase more strongly with repetition. This relationship was the most pronounced when taking all repetitions into account, and apparently absent for early repetitions. As discussed above, this may be related to two opposing processes affecting early repetitions, whereas effects for later repetitions are dominated by the effect that tends to increase gamma-band responses. The correlation between response strength and increase with repetition is in line with Brunet et al. (2014). They reported that for the same stimulus, the most responsive sites showed the strongest increases with repetition. We extend this finding, demonstrating that the strongest increases tend to result from stimuli that drive a site most strongly. Furthermore, the effect generalizes to a larger class of visual stimuli. Importantly, Brunet et al. (2014) also report that for putative pyramidal cells specifically, the increase in spike-field locking with repetition was strongest for the most strongly driven cells. This finding is difficult to replicate in the current dataset, since only MUA was available and there were relatively few trials to estimate changes in spike-field locking. Assuming that the effect reported by Brunet et al. (2014) generalizes, it would imply a stronger synchronization-based sharpening of responses to stimuli that generate strong gamma-band response to begin with. In future experiments, it would be interesting to contrast changes in perception with repetition for stimuli that do or do not generate strong gamma-band responses.

2.5.7 Advantages and limitations of the stimulus set

We presented colorful images of isolated objects. The objects were similar to those naturally encountered and relevant in the animals' daily environment (vegetables, fruit and brightly colored toys). We reasoned that this choice of stimuli might provide an engaging task for the animals and ensure that the visual input was not at odds with the natural image statistics their brains deal with normally (Fiser et al., 2010; Berkes et al., 2011). Importantly and in stark contrast to Brunet et al. (2014), next to being colored, naturalistic photographs, the stimuli were also novel for the animals on the first day of presentation. This way, it was possible to investigate the nature of repetition effects for images that are not overtrained. Overtrained stimuli may have considerably restructured visual cortex (Schoups et al., 2001; Woloszyn and Sheinberg, 2012; Cooke et al., 2015), such that a generalization of the effect to novel stimuli is not trivial.

The images were chosen such that they were likely to be perceived as a single unified object by the animals according to Gestalt laws. For example, a blossom of a rose was included, but not a potted flower. The latter may be perceived as either one object, or separated into the pot, the flower and possibly even the different leaves and blossom. This selection of images should therefore prevent or reduce possible rapid alterations in the monkeys attentional allocation that can occur naturally when several objects are present in the visual scene (e.g. Re et al. 2019; Landau and Fries 2012), which may have unknown influences on potential stimulus repetition effects.

Stimuli were large and typically overlapped with the fixation spot as well as larger parts of the lower right visual field. This enabled spatial uncertainty for the change detection task, and covered the RFs as well as their surround. Surround stimulation typically decreases MUA responses (see previous sections) and likely contributed to the relatively strong gamma-band responses observed for these stimuli (Gieselmann and Thiele 2008 and Peter et al. 2019, which will be discussed at length in Chapter 4). The highly colorful nature of the stimuli, intended to make stimuli easily discriminable, very likely also contributed to strong gamma-band responses and relatively weak rate responses (see Chapter 5, Rols et al. 2001; Shirhatti and Ray 2018; Peter et al. 2019). As a consequence, the dependence of repetition effects on response strength observed here may result from the particular class of stimuli studied. In particular, large, colorful, and high-contrast stimuli appear ideal to generate strong gamma-band responses, and as a consequence, strong increases in the gamma-band with repetition (see also General Discussion 6.2). Interestingly, this appears to contrast with effects for MUA responses. Large stimuli induce very little adaptation in V1 or LGN spiking, in particular when their contrast is high (Camp et al. 2009; Wissig and Kohn 2012, and citations therein). Both size- and contrast-dependence of adaptation speak to an intricate link with surround suppression. It has even been suggested that surround effects serve to protect responses from adaptation (Camp et al., 2009). Repetition effects in gamma-band activity may therefore be most pronounced when MUA responses change the least, in terms of stimulus space. Observing gamma-band synchronization can therefore tap into mechanisms that may be very subtle in terms of MUA responses. These observations put gamma-band activity at the crossroads of spatial and temporal context (surround suppression and repetition effects/adaptation) and their interaction. The interaction of spatial and temporal context may be central to vision (e.g. Schwartz et al. 2007). This topic will be explored further in the General Discussion and Chapters 4 and 5.

2.5.8 Role of attention, reward and behavior

The present experiment engaged the animals in a change detection task. Since stimuli were initially novel to the animals, task difficulty could not be equated between stimuli. Even if that had been achieved, it would be difficult to demonstrate conclusively that the animals perceived stimuli as equally relevant and that top-down processing was equal. Regression analyses that took reaction times and pupil responses (as indicators of stimulus-specific difficulty or arousal) into account still showed significant, stimulus-specific repetition effects for both MUA and gamma-band responses. Regression analyses focused on behavior showed no dependence of reaction time on repetition. Furthermore, the observed repetition effects were limited to MUA responses and gamma-band (or higher) frequencies, but did not affect lower frequencies. This is an important distinction, since lower frequencies have been related to top-down feedback effects, whereas higher frequencies have been related to feedforward processing

(Bastos et al., 2015; Michalareas et al., 2016). We further note that gamma-band responses for a given stimulus could first decrease, and then increase, whereas changes in attention with stimulus repetition are presumably unidirectional. Taken together, these results speak against a purely feedback (or top-down, attentional) account of the observed effects. A purely top-down, attentional effect may also show position invariance (see next chapter).

For V1 rate responses, adaptation or repetition effects are frequently studied under anesthesia (Solomon and Kohn, 2014; Kohn, 2007; Wissig and Kohn, 2012). In IT firing rates, no task-dependence was observed comparing passive viewing and change detection (De Baene and Vogels, 2010). Notably, rate increases can be observed when the repetition itself becomes task-relevant, and a repeated target stimulus has to be discriminated from other stimuli that also occur repeatedly (Miller et al., 1993; Miller and Desimone, 1994; Vogels and Orban, 1994; Eskandar et al., 1992). Vogels (2016) notes that differences in attention with respect to repetition resemble a chicken-and-egg problem: a decrease in response may reflect decreased saliency, or decreased saliency may result in response changes. In our data, we note that repetition effects occur rapidly, excluding an interpretation in terms of attention of these effects.

Finally, as is common in awake monkey electrophysiology, successful task performance was rewarded (in this case with fruit juice). On the one hand, even passive, non-rewarded visual stimulation can induce plasticity in awake rodents and insects (Cooke and Bear, 2010; Stopfer and Laurent, 1999), and stimulus repetition effects exist in primates under anesthesia (Solomon and Kohn, 2014; Kohn, 2007; Wissig and Kohn, 2012). On the other hand, reward (or the expectation thereof) can gate plasticity (Roelfsema et al., 2010), especially for long-term learning, such that its role in the current paradigm is unclear.

2.5.9 Current limitations and outlook

A number of limitations of the present study have been named throughout the preceding sections of the discussion. In particular, the precise stimuli and experimental design chosen show a generalization of the stimulus repetition effect to more naturalistic conditions. However, large and colorful stimuli are only a specific subset of natural stimuli. Furthermore, stimulus durations and other aspects such as prolonged fixation, and the pairing of visual stimulation with reward, limit the generalization of the observed effects to natural vision.

The chronic array recordings employed here allowed the current experimental design, comparing the effects of stimulus repetition for a given stimulus across days, with varying neighbors in time. However, single-unit recordings could provide important additional insights into the phenomenon. For example, Brunet et al. (2014) showed that putative interneurons play a fundamental role in increasing gamma-band synchrony with stimulus repetitions. What is the role of interneurons versus excitatory cells for the initial few repetitions? Furthermore, if a single cell participates in gamma-band synchronous encoding of several stimuli, will it show increasing locking selectively for stimuli that it "prefers" in terms of firing rate? Could a single cell participate in several gamma-defined ensembles depending on the stimulus, and change its locking behavior accordingly?

The simultaneous multi-unit recordings from arrays have the advantage that not only changes in single sites, but coordinated changes in the population of recording sites can be analysed. Preliminary analyses using Naive Bayes decoding suggest that later repetitions are more self-similar both in MUA and gamma-band responses, such that

training on later repetitions shows better transfer across repetitions, whereas responses to the first presentation are more poorly decodable from other presentations (data not shown). However, MUA responses for initial presentations are stronger overall, and decoding stimulus identity among only first presentations was not weaker than decoding identity from only one specific later presentation. This illustrates the difficulty of assessing changes with repetition using population decoding, since such measures can be affected by signal-to-noise ratio. Kaliukhovich et al. (2013) report that a single, immediate repetition reduces classification accuracy of the repeated stimulus in a small population of IT neurons. However, repeating one stimulus could enhance classification of another stimulus. The possible functional benefit of repetition will likely depend on the frequency of repeated versus novel stimulus encounters in natural vision, and the amount of cross-adaptation between these stimuli. Given the complexities of stimulus repetition effects, including several mechanisms with different time constants, it is likely that repetition effects on population decoding measures depend on stimulus choice and experimental design.

Furthermore, the current design revealed some long-term memory in the fixation behavior of the animals. It will be very interesting to see if changes observed within a session translate to some form of long-term plasticity. Since the present design maximized differences in stimulation between days, it is not designed to be sensitive to detect long-lasting response changes. A test for long-term effects should also include control stimuli that are not repeatedly presented across days.

We found that gamma-band increases with repetition depended on initial response strength. Given the strong stimulus-dependence of gamma-band activity, this suggests that gamma-band increases are particularly pronounced in a regime that involves large, high-contrast stimuli linked to surround suppression. The link between surround suppression and gamma-band activity will be explored further in Chapter 4. Such a link would predict that repetition effects in the gamma-band will be strongly dependent on stimulus size, in a way that may be inversely related to the size-dependence of repetition effects in MUA responses. Furthermore, it suggests that, for example, a weakening of the RF surround could weaken gamma-band repetition effects.

Finally, note that the overall strengthening of gamma-band responses for large stimuli is not a balanced plasticity process, with some sites showing decreases and other sites showing decreases in responses. This suggests that the effect for a given stimulus may asymptote to a ceiling, and may require homeostatic mechanisms such as sleep to restore responses back to a lower level. In the following chapter, these aspects will be considered when testing for stimulus position specificity and memory of the gamma-band repetition effect.

2.6 Chapter summary

1. Large, colorful, naturalistic stimuli induced a wide range of MUA and gamma-band responses in macaque V1.
2. Both MUA and gamma-band repetition effects were stimulus specific. No significant repetition effects were observed in lower frequency bands.
3. Stimulus repetition effects in the gamma-band generalize to naturalistic and previously novel images.
4. MUA repetition effects were characterized by decreases in particular for the first repetition(s). Reliable response reductions could be observed already in the initial transient response (< 100 ms).
5. Gamma-band activity could decrease or increase for the first few repetitions, and typically showed increases for later repetitions. This indicates that stimulus repetition induced multiple plasticity effects, operating at different time scales.
6. The dependence of gamma-band increases on initial response strength, and the stimulus-dependence of gamma-band activity, suggest that gamma-band increases are particularly pronounced in a regime that involves large, high-contrast stimuli linked to surround suppression.

Chapter 3

Location specificity and persistence of repetition effects

3.1 Acknowlegdements for Chapter 3

The data in this chapter was obtained from two animals originally trained and implanted (in cooperation with Richard Saunders) in the lab of Michael Schmid (see also previous chapter), that could be used for these experiments based on a new permit. I am particularly indebted to Dr. Michael Schmid, who enabled these experiments through a generous long-term loan of recording equipment.

3.2 Introduction

3.2.1 Location specificity as a signature of lower visual areas

Preliminary remarks: This chapter focuses on gamma-band responses, excluding a discussion of MUA. One reason is that in one animal, there was a sudden, chronic loss in signal from the V1 array. Given that in this animal, there were no reliable differences between average MUA responses when the RFs were stimulated compared to when they were not, the MUA data was excluded. In the second animal, there were clear MUA differences between the stimulated and non-stimulated condition on average. However, single-trial fluctuations were so high that the distributions of stimulated and non-stimulated responses overlapped (in contrast to single-trial gamma power estimates, as illustrated below). The likely reason for this is that a large stimulus was used, inducing surround suppression and consequently very low response magnitudes. See Chapter 4 for an extensive discussion of the relationship of surround suppression to gamma-band responses and rate responses. Further note that changes in gamma-band power are not displayed on log-scales, because in this design, stimuli induce more similar amounts of average gamma power.

In the previous chapter, it was shown that stimulus repetition affects both V1 firing rates and synchronization, in particular in the gamma-band, and in a manner specific to the stimulus. In theory, these effects could originate either 1) within V1, or 2) derive from upstream inputs (LGN or retina) or 3) result from feedback from downstream areas. Distinguishing between the first two possibilities is beyond the scope of the current study (see also Discussion). The possibility of feedback as a source of the effect will be addressed using recordings in V1, making use of the fact that visual areas high in the hierarchy have very large receptive fields that cover large parts of the visual field (see General Introduction, Figure 1.1). An effect that originates in area IT or a similarly higher area can therefore show invariance to a change in stimulus position (Sawamura et al., 2006; De Baene and Vogels, 2010). Conversely, an effect

that is sensitive to changes in position points to the involvement of lower to mid-level areas with smaller receptive fields. Adaptation in one part of the RF did not transfer to another part of the RF, suggesting that the effect was inherited from an earlier area. The following experiment will therefore test for a transfer of the repetition effect across stimulus location.

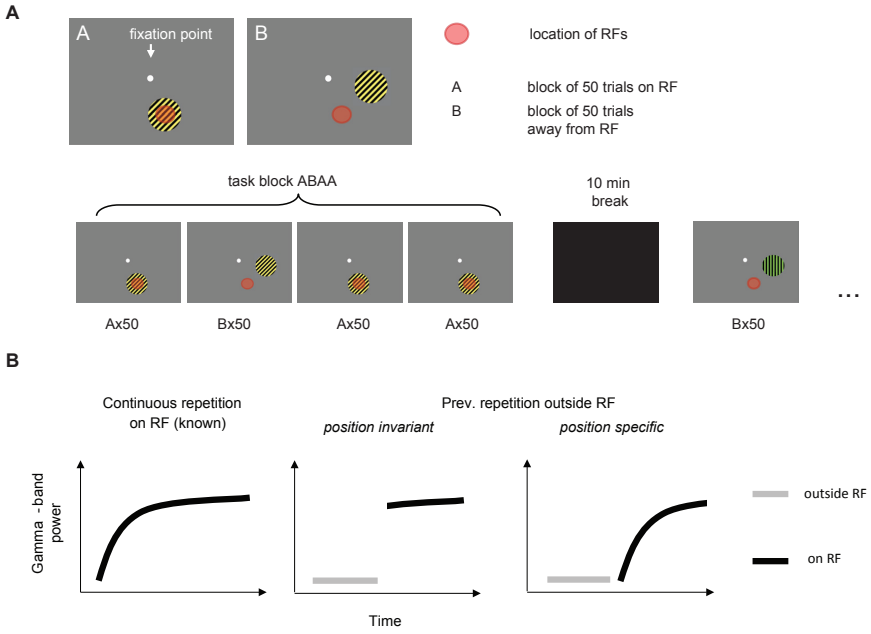

Figure 3.1 — (A) Illustration of the task. Top, positioning of the stimuli on the monitor and approximate RF locations of the recording sites of the array (red). Stimuli were shown in blocks of 50 trials either on the RFs (condition A) or away from them at the same distance from the fixation point (condition B). Bottom: Illustration of the task sequence of one day. One task block consisted of 4x50 trials, e.g. ABAA, followed by a break, another task block, and so on. (B) Illustration of alternative hypotheses. Based on the known fall-off of the rate of repetition increases in gamma-power over time with more reptitions (left), a location invariant mechanism would predict that after repeated stimulation at another location, the rate of repetition increase is low. The slope of the increase should depend on the total number of repetitions. In contrast, a location specific mechanism would predict a higher rate of increase when the local number of repetitions is lower than the total number of repetitions.

Specifically, a single, large grating stimulus was repeatedly presented in a blocked fashion at two alternating locations, with blocks in specific sequences (see Figure 3.1A). A large stimulus (4.25 dva radius) was chosen in order to cover all receptive fields that were recorded from the chronic arrays (see Methods of previous Chapter 2) as well as the receptive field surrounds. Coverage of RF surrounds is a necessary condition to generate strong gamma oscillations in V1 (both in case of grating stimuli and otherwise, see Chapter 4). Block presentation was used to make use of a particular property of the stimulus repetition effect: its dependence on the repetition number in a log-linear way. In other words, repetition effects are strongest for the initial presentations, and then taper off to ever-smaller increases or steady responses (see also Discussion). Consequently, fitting a slope to gamma-band responses with repetition will yield a large slope at the beginning and a smaller slope with later repetitions (see Figure 3.1B for an illustration, Figure 3.7 for example of effect of rest). Therefore,

concrete predictions about the rate of gamma-band increases with repetition can be made, based on the hypothesis that repetition effects are originating in lower visual areas and hence are location specific. Namely, assuming location specificity, the slope of the gamma-band repetition effect should be unaffected by previous stimulation at a different location (Figure 3.1). The location specific, rather than total repetition number, should determine the slope of the increase. This allows a comparison of the slope of the repetition effect given previous stimulation at either the same, or a different location. This equates passage of time, the task, and the number of experienced rewards, manipulating only the stimulation history. Given location specificity, the slope in a later stage of the experiment should be lower if the previous stimulation was at the same location, but high if this is the first time this location was stimulated.

3.2.2 Persistence of repetition effects

Another open question regards the duration of the effects of stimulus repetition. In Chapter 2, we saw that repetition effects can build up inspite of intervening stimuli on the timescale of seconds. This in itself could be considered a remarkable feat for V1, which is often thought to integrate information only on sub-second timescales (Ringach and Shapley 2004; Kim et al. 2019, see Discussion). It has also been shown that a break in the task of several minutes, during which the animal rests, can induce a reset in the repetition effect, such that gamma-band responses decrease to the level of the beginning of the experiment (see Brunet et al. 2014 for a break initiated by the animal, see Figure 3.7 for an example of planned breaks as part of the experiment). Similarly, some reset of the repetition effect most likely occurs over night, which arguably involves a much longer rest. Here, we ask whether the repetition effect can last at least on the timescale of minutes, rather than seconds, making use of the location specificity paradigm (Figure 3.1). Given a blocked presentation of a stimulus at two alternating locations, and a location specificity of the effect, one can ask what happens after a return to a stimulus location after a "break" in the form of stimulation at the other location for several minutes. Again making use of the change of slope of the repetition effect, one can test if the stimulus history several minutes back still matters. Alternatively, there may be a full reset to the initial gamma-band response strength similar to a resting break of a few minutes (see Figure 3.2 for illustration).

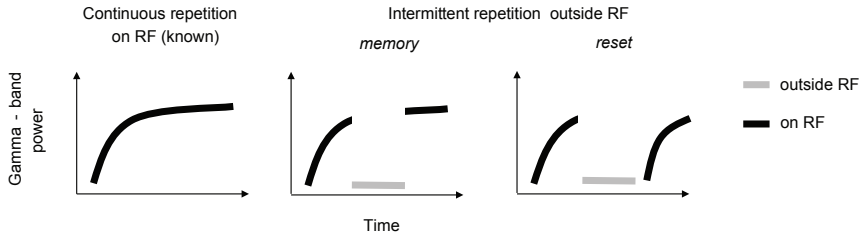

Figure 3.2 — Illustration of alternative hypotheses for a memory effect. Based on the known fall-off of the rate of repetition increases in gamma-power over time (left), a mechanism involving memory inspite of prolonged stimulation at a location away from the RFs would predict that gamma-band power remains at the level achieved by previous location-specific stimulation, rather than reducing to the level at the beginning of location-specific stimulation.

This question is interesting because rest and sleep have been proposed to enable homeostatic mechanisms that concurrently allow a refinement of sensory responses (Vyazovskiy et al., 2008; Huber et al., 2013; Kuhn et al., 2016). Rest could involve a homeostatic normalization of overall excitability that respects recently potentiated connections, and thereby "convert" current changes on the functional level, such as strong gamma responses, to a structural change. This would require a "memory" of recent stimulus repetitions, possibly leaving a trace in functional responses as e.g. in the form of stronger gamma responses, until such rest can be experienced (see Discussion). Consequently, a survival of a repetition effect on the timescale of minutes, even if not visible across days, could serve an important function.

3.3 Methods

3.3.1 Task

Recordings were obtained in monkeys H and K as described in the previous chapter. The monkeys performed a change detection task on colored, square-wave grating stimuli (static, spatial frequency 2 cycles per dva, radius 4.25 dva). The stimulus was either centered on the V1 receptive field locations in the lower right visual quadrant, or on an equi-eccentric location near the horizontal meridian. In monkey H, the fixation spot was moved up 2 dva from the monitor center, since receptive fields were relatively eccentric and this allowed the placement of the stimulus on the RFs. Monkeys maintained fixation on a white, circular fixation spot 0.2 dva in size, for 1.3 s of gray background stimulation followed by 1.5-2.3 s of grating stimulation. In this variable interval, a circular stimulus change of 0.4 dva diameter could occur at a random location on the stimulus. The random locations were restricted such that the full changed spot remained within the stimulus. Monkeys reported the change by an eye movement toward the change location within maximally 1 s to obtain a juice reward. Task difficulty was maintained at a level that was not challenging to the animals. The change detection task only had the purpose of keeping the animals engaged and was not in itself of interest. From extended experience with the animals, it was clear that a more challenging task will lead to an increase in the variability of inter-trial intervals over the session. As the animal becomes more sated or less aroused over time, self-initiation of trials becomes more variable. Additionally, more and more errors are made if the task is difficult, inducing changes in block lengths. This variability would be reflected in the neuronal data. As an additional measure to maintain a steady engagement during stimulation blocks, a ten minute break with a dark monitor followed every 200 trials (a "task block") of stimulation. This encouraged steady responses during task blocks. A further purpose of this break was to use the known reset effect of such breaks on gamma-band responses (see also Figure 3.7). This enabled more data collection on the same recording day. In an additional attempt to ensure independent data in different task blocks, the stimulus was altered on every block. Specifically, a cyan grating stimulus of 45 degree orientation, a green vertical stimulus and a yellow 60 degree oriented stimulus was used. Note that these stimuli were not equally bright, nor did they have equal luminance contrast. Response differences between stimuli are not of interest in this experiment, and all comparisons are made within-stimulus. All three stimuli elicited robust gamma-band responses. The stimulus order in the blocks was fixed across days, in order to compare responses to the same stimulus in the same part of the session across days. The central manipulation of the experiment was that within each task block, stimulus position could change every 50 correct trials

("miniblock"). The design was counterbalanced, such that the overall probability of stimulation occurring at one location or the other was 50%. With locations A and B, A being on the V1 receptive fields, 8 possible sequences were used in a task block: all A, AABB, ABAA, or BAAA (4 types of A-Blocks), and their inverse (all B,... 4 types of B-Blocks). On a day, A-Blocks and B-blocks alternated, and the starting block alternated between days. The sequence type was assigned in a pseudorandom manner.

In a single session in monkey H, it was investigated whether the repetition increase would show memory after "interference" with rapid, repeated stimulation with stimuli of different orientations. An achromatic, maximal contrast, moving square wave grating (1.0 dva spatial frequency, 0.5 dva temporal frequency, 7 dva diameter, position centered on V1 receptive fields) was shown in a passive fixation task. A stimulus of vertical orientation was repeated 150 times, followed by the "interference" sequence, and another 150 trials of the vertical stimulus to test for disruption of the repetition effect. After this, a 10 minute break followed, and a new repetition-interference-repetition sequence started. The interference block consisted of trials with either other orientations (30, 60, 90, 120, 150 degrees angle, first and last interference block in the session) or of trials including those other but also the repeated orientation (middle interference block). Each trial in the interference block consisted of 200 ms individual grating presentations with 100 ms interstimulus interval, equating the total trial duration to the trial duration of the repeated stimulus (1.4 s, baseline duration 1.3-1.4 s). Individual interference trials contained a randomly permuted sequence of orientations. The interference block contained 25 trials, or 125 stimulus presentations.

3.3.2 Data analysis

Preprocessing and trial selection. For both animals, these recordings were made using the implants described in the previous chapter. However, instead of recording using the CerebusTM Neural Signal Processor, data was acquired using Tucker Davis Technologies (TDT) systems using an adapter from the connector to the TDT PZ2 preamplifiers. MUA was estimated by band-pass filtering (300 Hz-12000 kHz) with a 4th order zero-pass Butterworth filter, and filtering and downsampling to 1/24th of the original sampling rate using an 8th order FIR filter.

LFP signals were rereferenced through subtraction of the mean across recording sites. Single-trial average gamma power estimates were obtained as described in the previous chapter, using single-trial (instead of a common) baseline correction, and a window of ± 8 Hz around the cross-site gamma peak (simple maximum across frequencies) per stimulus. Power was expressed as fold change from baseline (stimulus/baseline), "relative gamma power". Since stimuli were chosen to induce similar gamma-band power, the data is no longer plotted in a log-transformed fashion as in the previous chapter (the log-transform de-emphasizes changes with repetition).

Because of the possibility of a reset of repetition effects when the animal took a break, interruptions in performance of more than one minute led to the exclusion of all further miniblocks of a task block, and the exclusion of the current miniblock if less than 80% of it was completed (for the purpose of slope fitting). Additionally, the very first presentations of a task block, as well as location changes, could induce a transient decrease in gamma-band (and broad-band) activity (see Figure 3.4, see also the previous chapter). The first 10 trials of a block were therefore excluded from analysis (in monkey H, similar results were obtained when not excluding any trials, because the initial decrease was more pronounced in monkey K, see Figure 3.4).

Slope fitting and statistics. For the analyses in Figure 3.5 and 3.6, the 5 sites per animal with the strongest gamma-responses were selected and averaged. Similar conclusions (with significant results) were reached when slopes were evaluated for the cross-site grand average, and also in case of a bipolar derivation of the LFP signals (with bipolar pairs chosen from neighbors of different depths). For each 50 trial block, a simple first-order fit of the data was computed, yielding a slope and an offset. Since the design required comparisons across recording days, and the overall mean gamma-band activity induced by a stimulus can vary between days, a comparison of slopes was the focus of most analyses. Only slopes of the same time in the session and the same stimulus were compared. Significance of differences of slopes was assessed with a Welch's t-test, i.e. an unpaired t-test that does not assume equality of variances between distributions. On some occasions for slope comparisons, there were a few more slopes of one type than the one it was being compared to. In such cases, the slopes closest to each other in time are plotted. All slopes were used to compute the unpaired t-tests. Conclusions based on the plotted slopes only were equally significant.

Regression analyses were performed using Matlab's functions fitlm or fitlme, followed by an ANOVA to identify significant predictors, and a theoretical likelihood ratio test to compare different models. For regression analyses, only trials with the stimulus position covering the RFs were modeled. This was done because gamma power is trivially absent in the other case, and including the regression variable "location" would mean that other terms of interest would need to be modeled as interactions. Regression models predicting single-trial gamma-power were constructed using the predictors "monkey identity", "session", "stimulus identity", "reaction time", "eye position eccentricity", "eye position variance", and "pupil response" (response during the time that LFP power was estimated, divided by baseline pupil response), and the main predictors to test the hypotheses, the "log(total repetition number)" versus the "log(local repetition number)". The log transform was used because it yielded better fits, in line with previous findings by Brunet et al. (2014). Similarly to the slope fitting procedure, the first 10 trials of a block were excluded from analysis.

3.4 Results

3.4.1 Location specificity of repetition effects

Grating stimuli were presented at two possible locations, one of which overlapped with the RFs. Consequently, clear gamma-band responses were induced only in case of RF stimulation (Figure 3.3). 24 session were recorded in monkey H and 18 sessions in monkey K, each session typically contained at least 3 task blocks (i.e. 600 trials).

Stimulus repetition led to clear increases in gamma-band responses (Figure 3.4). In one animal (monkey K), the initial ca 10 repetitions induced a transient decrease in the gamma-band response, along with a downward shift in peak frequency. The downward shift in frequency is also observable in the second animal. As illustrated in Figure 3.4A, the effect was broad in frequency, but dominated by gamma-band activity similar to that induced by the stimuli later. Conclusions were similar regardless of the alignment of the trial axis to correct or all trials (including fixation breaks).

This means that the lack of the decrease in monkey H is not due to a mixing of trials at different parts of the session, and that the time to a switch to an increase truly was on the order of ca 10 repetitions for monkey K. Visual inspection of pupil responses and eye position variance developments with repetitions indicated that the transient decrease can be present without a change in pupil or eye position variance, although

the pupil showed stronger constriction for initial presentations of some stimuli, and for 1-3 presentations only (data not shown). The transient decrease most likely reflects a separate phenomenon from the typical increase in gamma-band responses with repetition (see also previous chapter). For the following slope and regression analyses, the first 10 trials of each block were therefore excluded.

To test for location specificity, the slope of the repetition effect was compared for the *second* (mini)block in each sequence, where the *first* block could have consisted of stimulus repetitions at the RF location (total repetition number = local repetition number) or of stimulus repetition outside the RF location (total repetition number > local repetition number). As illustrated in Figure 3.1B, a location-specific effect suggests that the slopes in the latter case will be higher - as only the local repetition number determines the rate of the increase. As illustrated by an example in Figure 3.5A, and then evaluated for all session pairs, the repetition effect indeed showed location specificity. The slopes were significantly stronger when the second block was the first in which stimulation occured at this location ($P < 0.05$). A similarly strong increase could also be observed if the third (rather than second) block in a task block was the first with stimulus presentation in the RFs (sequence BBAA, see next section on memory below).

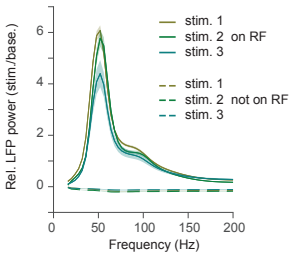

Figure 3.3 — Average LFP power spectra in response to visual stimulation with the three different types of grating stimuli for monkey H. Clear gamma-band responses occur only when the stimulus is on the RFs.

The location specificity of the repetition effect was corroborated by a regression analysis. A regression analysis predicting relative gamma-band power in case of stimulation of the RFs (see Methods) showed that a model that included both local and total repetition number significantly outperformed a model that included only total repetition number[1] (theoretical likelihood ratio test, $P < 0.001$). Furthermore, a model that included only local repetition number significantly outperformed a model that included only total repetition number ($P < 0.001$). However, both local and total repetition number were found to have a significantly predictive effect on gamma-band power, with both repetition effects explaining similar amounts of variance on average (log(local repetition number) beta = 69.36, r-squared = 0.006, log(total repetition number), beta = 65.48, r-squared = 0.005, both $P < 0.001$). When restricting the analysis to the miniblocks that were explicitly compared during the slope fitting, (i.e. data which maximally decorrelates total versus local repetition number), effects became more pronounced. Pooled across animals, there was no longer an effect of total trial number ($P = 0.53$), and the amount of variance explained by local repetition number increased (beta = 164.68, $P < 0.001$, r-squared = 0.047). The regression models included further control predictors for “pupil responses”, “eccentricity”, “variance of the eye position”, and “reaction time” (the available measures of overt behavior), as well as “stimulus identity” and “monkey identity”. Apart from stimulus and monkey

[1] In both total and local cases, the predictor was log(repetition number). Note that this repetition number here refers to the number *in the task block*. Given the 10 minute break and switch of stimulus identity between task blocks, the local repetition number across the entire session should play little or no role. Indeed, a regression model including the total versus local number of repetitions *across the entire session* showed no significant effect of local stimulus repetition, $P > 0.3$.

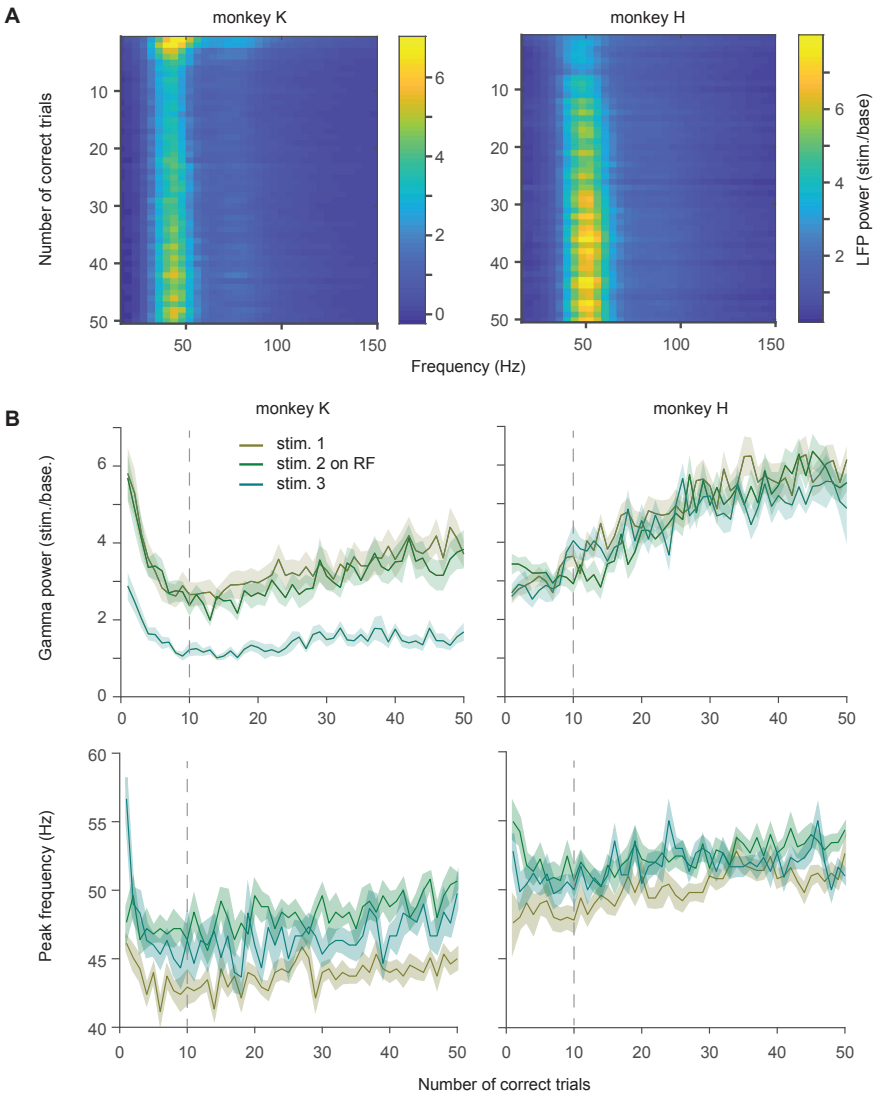

Figure 3.4 — Analysis of repetition effects in a miniblock. (A) Cross-session average spectra of the first 50 trials of RF stimulation of each task block for stimulus 1. Left panel: monkey K. Right panel: monkey H. Note the transiently strong responses in monkey K, and the upward shift in peak frequency with repetition. (B) Top panels: Cross-session average gamma-band power (averaged around peak of each stimulus, see Methods) of the first 50 trials of RF stimulation for all three stimuli. An initial decrease is clearly visible in monkey K. Bottom panels: Cross-session average peak frequency (in the range 25-80 Hz) of the first 50 trials of RF stimulation for all three stimuli. Gray dashed lines indicate cutoff for slope fitting analyses. Fitting slopes to the increase in peak frequency with repetition after trial 10 yielded significantly positive slopes in 5 out of 6 stimuli in the two animals based on a permutation test and multiple comparison correction.

predictors, these predictors typically captured less variance then the predictors of interest, and were not always significant in individual animals. The conclusions about the local vs. total repetition number held both with and without inclusion of these control predictors in the model, indicating that the repetition effects are not a spurious effect of changes in behavior. Taken together, these results indicate that gamma-band repetition effects show location specificity. The regression analyses indicate that some additional effect of the total repetition number may be present (see Discussion).

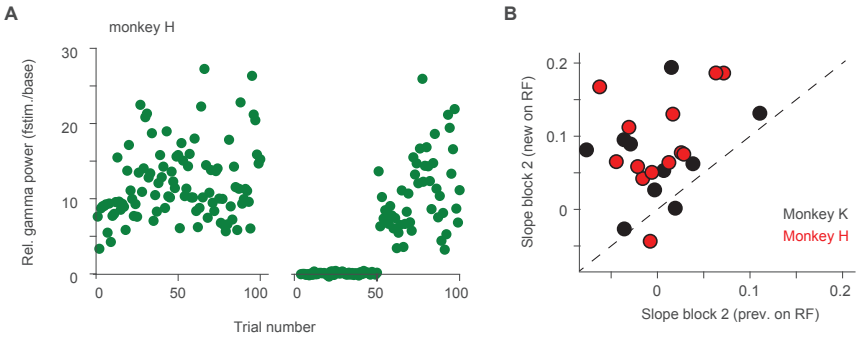

Figure 3.5 — Test for position specificity. (A) The first 100 trials from two example task blocks from two separate recording days. Left panel: example of continuous repetitive stimulation on the position covering the RFs (sequence AA). Right panel: example of the first 50 presentations away from the RF locations, followed by repetitions on the RF locations (sequence BA, see Figure 3.1). (B) Comparison of slopes for the second block of a sequence of continuous repetitions on the RF (A**A**), versus slopes for the second block, where this is the first block with stimulation on the RF (B**A**). Slopes were significantly stronger in case B**A** ($P < 0.05$).

3.4.2 Persistence of repetition effects

Having established location specificity, the same paradigm could be used to test for memory effects of stimulus repetition. Since the repetition increase on gamma-band power was location specific, a repeated stimulation of the RFs, interrupted by a block of repeated stimulation away from the RF, effectively constituted a break of several (ca. 5) minutes. A break that is self-initiated by the animal and constitutes some form of rest has been shown to reset gamma-band responses to levels similar to initial repetitions (Brunet et al., 2014). Here, it is tested what happens when the animal continues to be engaged in a visual task, which includes a maintained level of arousal and reward, but an interruption of the local visual stimulation. Specifically, the possibility of a reset of the gamma-band response, resulting in a stronger rise in gamma-band responses with repetition, is tested against a memory effect of the gamma-band response and therefore maintained gamma-band response with little further increase with repetition (see Figure 3.2 for an illustration). Note that this test requires that location specificity is established first, because if stimulus repetition showed transfer across positions, the change in location would not constitute a break. Lower slopes after a stimulation at another location would be expected in such a case due to transfer from the other location. To compare the situation of a return to a previously stimulated position (a sequence starting with ABA) with a situation of equated time in the session, experienced rewards and therefore arousal, the third block of a sequence without previous stimulation (BBA) was used. Figure 3.6 shows an example session

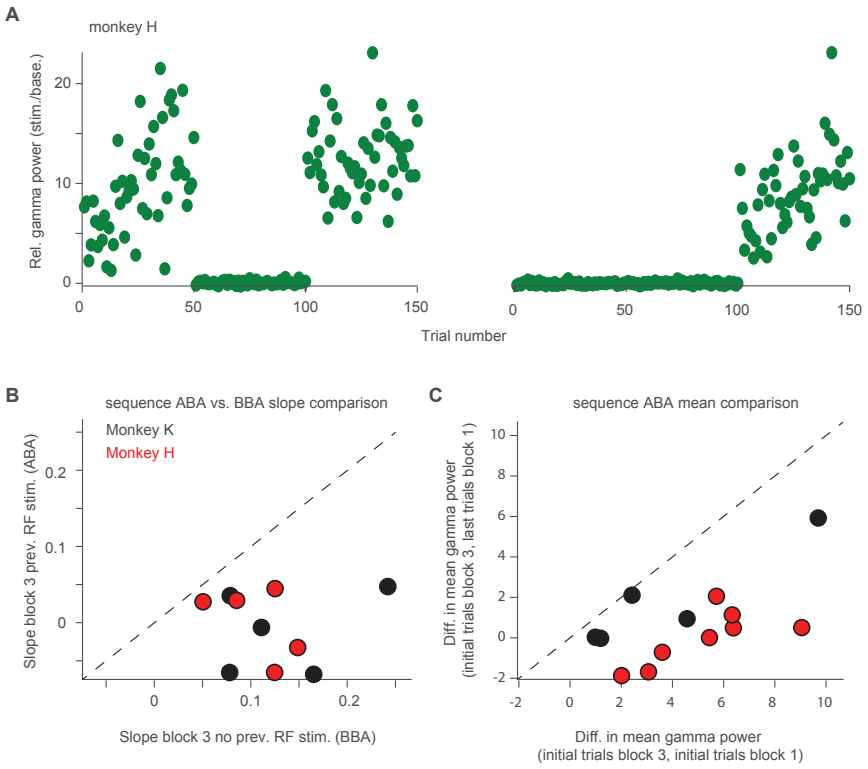

Figure 3.6 — Test for persistence of repetition effects. (A) The first 150 trials from two example task blocks from two separate recording days. Left panel: example of repetitive stimulation on the position covering the RFs, with an intermittent block at the other position (sequence ABA). Right panel: example of the first 100 presentations away from the RF locations, followed by repetitions on the RF locations (sequence BBA, see Figure 3.1). (B) Comparison of slopes for the third block of a sequence AB**A**, versus slopes for the third block, where this is the first block with stimulation on the RF (BB**A**)). Slopes are significantly higher in case BB**A** than AB**A**. (C) Comparison of differences in mean gamma power responses for the initial trials of the third block in sequence ABA, to either the initial trials or the last trials of the first block in the sequence. The initial trials of the third block are significantly more similar to the last trials than the first trials of the first block, indicative of a memory effect. Significance holds for monkey H or pooled data, data points for monkey K all go in the same direction, but not significantly so. Initial trials are defined as trials 5 to 10 (see also Methods), and last trials as trials 45 to 50 in a block of 50. Note that because these comparisons are made within-session, a comparison of mean responses is feasible. Also note that compared to the between-session comparisons in (B), this results in additional data points (see also Methods).

of type ABA, demonstrating a clear memory effect, and an example session of type BBA, showing location specificity (by the clear increase despite prolonged previous stimulation). Similar effects could be observed across the session pairs per animal, as illustrated in the scatter plot in Figure 3.6B, and were significant ($P < 0.05$). In contrast, slopes of sessions of type ABA and BAA, compared for the last A block, were not found to be significantly different in either animal (data not shown). While one should be careful not to overinterpret the absence of an effect (which could be due to lack of sensitivity), a reset hypothesis would suggest that a difference should be found in this situation, and this was therefore tested here.

A further complementary analysis is to compare the strength of gamma-band responses at the beginning of the third block in the ABA sequence (i.e. first responses after the return to the original position) with the beginning versus end of the initial block. Given a memory effect, the responses should be more similar to the end, rather than beginning, of the first block. Indeed, mean responses at the start of the third block are on average significantly more similar to the mean responses at the end, rather than beginning, of the first block (Figure 3.6C). Note that a comparison of mean responses is a sensitive analysis here because it is performed within a task block, rather than across days.

Finally, regression analyses were performed also for the memory effect. Namely, the same miniblocks (AB**A** and BB**A**) as for the slopes were analysed in the model. These represent cases where the total repetition number is the same, and the difference in local repetition number is resulting only from a miniblock not directly preceding the current block. Therefore, any significant effect of the local repetition number is indicative of a memory effect (because if we had started the counting of local repetitions up to 50 trials before the current block, it would be identical). As expected from the slope analyses, the regression models showed a significant effect of local repetition number. Similar to the regression analyses for the location specificity effect, there was an additional effect of the total trial number (both $P < 0.006$, local repetition number r-squared = 0.017, total repetition number r-squared = 0.006).

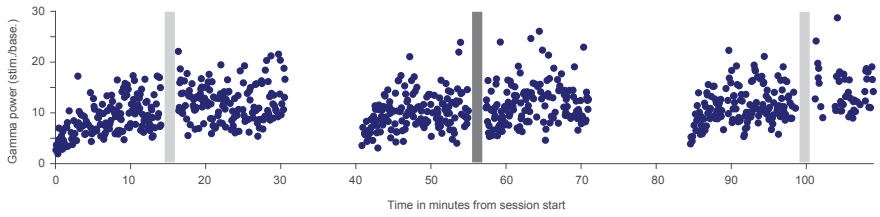

Figure 3.7 — Test of reset with stimulus interference and illustration of reset with breaks. Single session from monkey H testing a "reset" or disruption of the gamma-band increase with repetition through rapid presentation of other grating stimuli. 150 presentations of a moving, achromatic grating stimulus were followed by an "interference" block of either only other orientations (light gray) or also the repeated orientation (dark gray). After this, another 150 presentations of the repeated grating orientation followed. Visual inspection suggests that the interference blocks of 125 stimulus presentations in 25 trials (see Methods) did not result in a reset of gamma-band responses. In contrast, the breaks of 10 min in between the repetition-interference-repetetition sequences did result in a reset of gamma-band responses.

A memory effect for the gamma-band increase while the animal is performing a task at a different location is a powerful proof of memory, because no other stimulus at the location could have maintained the (then necessarily stimulus-unspecific) gamma-band response increase. On the other hand, the network, after repeated stimulation, is not exposed to patterned stimulation during the time when the stimulus is at the position outside the recorded RF position[2]. It could be that memory is maintained only in such a special situation. In one recording session in monkey H, this notion was tested by attempting to interfere with the repetition effect (Figure 3.7). After 150 initial repetitions of a single stimulus, rapid, massed repetition of stimuli with

[2]Note that the monkeys moved their eyes freely between trials, i.e. some form of patterned stimulation did occur.

different orientations from the repeated one followed for 25 trials (i.e. 150 stimulus presentations) or about 2 minutes. Another 150 presentations of the initially repeated stimulus followed. It was visible that the gamma-band responses maintained the level acquired during the initial repetition block inspite of the attempted interference. In contrast, a ten minute resting break did result in a disruption of the repetition effect (Figure 3.7).

3.5 Discussion

The above experiments used massed repetition of a simple grating stimulus to investigate the position specificity and memory effects of the increase of gamma-band responses with repetition. This design was used because the previously observed trajectory of the repetition increase, namely a stronger increase for the first presentations compared to a later, more stable period, allowed straightforward tests of the position and memory hypotheses, and effect sizes that allowed direct observation of the effects even during the recordings. Clear position specificity and memory for previous repetitions were observed. These effects will be discussed in turn.

3.5.1 Location specificity of repetition effects

Studies of stimulus repetition in high-level visual area IT have used the (partial) position invariance of repetition effects there to argue for an involvement of IT-level visual areas. The reasoning is that the much smaller size of receptive fields in lower visual areas prevent position invariance there (De Baene and Vogels, 2010; Sawamura et al., 2006). Here, the argument is inverted: any general top-down feedback effect impinging on V1 from higher-order visual areas should show position invariance. Position specificity is not expected in such a case. The above analyses show that the gamma-band repetition effect was strongly, though not exclusively, position specific. In particular, the regression analyses indicated that some additional effect of the total repetition number may be present. This general effect was not specific to the task blocks (and therefore stimuli), but even present throughout the entire session. It is therefore possible that this "global" repetition effect captures unspecific effects, such as aspects of arousal that are not captured by the pupil response changes (which can reflect, among other factors, arousal levels, Ebitz et al. 2014; Bradley et al. 2008)[3]. The current design can detect position specific effects, but cannot determine the origin of any further, non-specific effect.

Taken together, the present results indicate that the gamma-band increase with repetition shows a substantial degree of position specificity. This excludes higher-order visual areas as the sole source of the effect. Consequently, this increases the likelihood that the repetition effects have an origin in plasticity processes within V1. However, alternatives to this scenario are plasticity processes in structures that provide feedforward input to V1 (predominantly retina, LGN) or in mid-level visual areas like V2 or V4 that also provide feedback to V1. Indeed, Brunet et al. (2014) showed that not only V1 but also V4 showed increased gamma responses with stimulus repetition,

[3]Furthermore, the regression analyses should be interpreted with caution. Since the local and total repetition number only decorrelates across certain miniblocks, but not *within* miniblocks, regression models that include both total and local repetition number could split up variance between those factors in some way. Conversely, a model that only includes the total repetition number will invariably fit any existing local effects to some degree. It is therefore difficult to rule out or confirm definitively any modulatory effects of the total repetition number.

and that coherence between the two areas in the gamma-band was increased. These findings are compatible with a scenario where effects in V4 are driven by plasticity in V1, or the other way around. It is also not unlikely that both areas show coordinated plasticity. V1-V2 gamma-band synchrony has been demonstrated for these types of stimuli as well, although not in the context of stimulus repetition (Roberts et al., 2013). It is therefore possible that this phenomenon encompasses (or coordinates) lower to mid-level visual areas.

Could all these effects result from changes to the input in V1, possibly cascading down the visual system? Such a cascade could derive from two different changes in the input. Either the input is in itself gamma-rhythmic, or it changes in some way that is translated into gamma-rhythmic responses within V1. Evidence that speaks against the former possibility exists in an indirect form. Based on the current state of the literature, thalamic inputs to V1 or other sensory areas are not themselves gamma-rhythmic (Bastos et al., 2014; Bessaih et al., 2018) and therefore not directly entraining V1. The laminar distribution of stimulus-induced gamma-band responses in V1 away from the input layers (see General Introduction, 1.4, Lowet et al. 2017; Xing et al. 2012) also speaks against a rhythmic entrainment via the input structures.

Some more general changes in V1 input, such as changes in LGN firing rates or changes in the thalamocortical synapse also deserve consideration. One way to increase gamma-band responses is an increase in input drive, for example through increasing stimulus contrast (e.g. Börgers and Kopell 2005; Lowet et al. 2017; Ray and Maunsell 2010). Due to the complexity of the circuitry, there are several ways this could come about. It should be noted that in general, a system with several network and single-cell time constants (i.e. V1) can bring about changes in the balance between inhibition and excitation that in some ways could increase gamma-band responses in many ways.

The following will illustrate one particular scenario as an example. Repeated stimulation tends to lead to steady or decreasing responses in the thalamus (Camp et al., 2009; Dhruv et al., 2011; Solomon et al., 2004). Intuitively, one would assume that this would result in decreases in V1 responses to these inputs. However, one should note that thalamocortical synapses may be in a continuous state of synaptic depression from high spontaneous activity in the thalamic input (Boudreau, 2005; Stoelzel et al., 2015; Castro-Alamancos, 2002). This would constitute a case for "less is more": when thalamic cells decrease their response rate, the effect of any given spike may increase, because the synapse is released from depression. Indeed, Stoelzel et al. 2015 have shown evidence for such an effect, using prolonged stimulation with a (small) moving grating that was optimized to maximally drive single LGN cells in the rabbit. LGN firing rates were reduced, but the layer 4 LFP response evoked by an individual spike was increased. One could imagine that fewer, but more effective LGN spikes (i.e. with greater depolarization in the receiving cells) could drive layer 4 (i.e. the input layer) of V1 more strongly, yielding a "paradoxical" rate increase. This would require that the increase in efficacy of a given thalamic spike outweighs the overall reduction in spikes. Assuming a generalization of this effect to large, surround-suppression inducing stimuli as used here, and a bias in our recordings as well as those of Brunet et al. (2014) to upper layers, it could be that the hypothetical rate increase is not detectable, yet driving gamma-band increases. The likelihood of this scenario is difficult to estimate, because adaptation effects in V1 for similar stimuli depend in a strong and qualitative way on stimulus size (see also previous chapter). This scenario was discussed in detail here to acknowledge the existence of adaptation effects on timescales studied here in V1 inputs, and illustrate that their effects on V1 may not be intuitive.

Finally, it should be noted that the LGN inputs to V1 receive strong feedback from V1, which has size-dependent effects (Sillito and Jones, 2002; Murphy et al., 1999; Andolina et al., 2012; Jones et al., 2012). Therefore, although V1 inputs such as LGN may change their responses with repetition, these effects may ultimately originate from V1 feedback. Adaptation effects in the retina and LGN are believed to be relatively stimulus-unspecific (Solomon and Kohn 2014 – but see Stoelzel et al. 2015). If this holds true, it would make a precortical origin of the repetition effect less likely, given its stimulus specificity.

An interesting test for a future study would be interocular transfer, i.e. whether gamma-band increases with repetition experienced in just one eye will be maintained if one switches to the other eye. This is because binocular cells to a large extent first appear in V1. Using such interocular tests, within-trial rate decreases with repetition show laminar signatures of a cortical origin in primates (Westerberg et al., 2019). This finding is also supported by optogenetic silencing of V1 in mice (King et al., 2016). Collectively, these considerations point to a retinothalamic or cortical locus of the repetition effect, anywhere from the first thalamocortical synapses to mid-level visual areas (and quite possibly on more than one level of the hierarchy).

Testing for location specificity of the repetition effect is a useful tool to constrain the mechanism behind the observed effect. Another question is under which circumstances location specificity is a useful property for any adaptive or plasticity process. After all, visual objects can occur anywhere in the visual field. The repetition effect may therefore reflect an adaptation to a temporally location-specific situation. For a person, this could be sitting at a desk, where parts of the scene will have a relatively constant position. Another possibility is that the massed repetition paradigm itself encourages a locally specific plasticity process. The blocked paradigm generates very strong predictions about the next stimulus position, which the visual system could use. To summarise, for the current massed repetition paradigm, the present results and current state of the field indicate that the most likely source of the repetition effects on gamma-band activity reside within V1 up to mid-level visual cortices. An inheritance of the effect from upstream input to V1 that is translated into a gamma-band increase cannot be excluded, however. Further insights into the origin of the effect could be gained by recording and manipulating other brain areas, and employing a greater variety of task designs, testing in particular for interocular transfer.

3.5.2 Persistence of repetition effects

The current experiment investigated the effects of prolonged repetition of a stimulus on visual gamma-band responses. It was shown that 50 repetitions of a stimulus (corresponding to a total of about 1.5 min of visual stimulation) can result in a memory effect over several minutes (> 5 min). This represents an understudied timescale, with many studies on repetition focusing on either very few repetitions, or long-term learning across days (see Introduction of previous chapter). In contrast to these types of paradigms that either are concerned with different timescales, actively prevent memory effects, or try to examine particular processes such as encoding or working memory, other paradigms explicitly use massed repetitions and look for repetition effects on an intermediate timescale. These typically concern themselves with either "adaptation", or simple forms of learning such as synaptic potentiation.

The observed transient decrease in gamma-band responses with initial repetitions (Figure 3.4) may reflect a simple form of adaptation due to, for example, input fatigue, and the later repetition increase (here: repetition effect) another form of adap-

tation due to synaptic changes or changes in inhibition-excitation balance. The size and persistence of repetition effects for massed repetitions of simple stimuli depends on many factors, chiefly among them the type of stimulus used and the amount of adaptation in form of prolonged and/or repreated presentation. The latter has been formalized in the so-called duration-scaling law, where the strength and persistence of repetition effects increases with the amount of repetition (or stimulus duration), but does not change in quality (Bao and Engel, 2012). Duration-scaling captures some strong effects in the adaptation literature (Bao and Engel, 2012; Solomon and Kohn, 2014). Given that gamma-band repetition effects are clearly cumulative, and that the current design employs relatively many repetitions, this means that the observed persistence of the gamma-band increase over minutes is likely the result of this particular design. Fewer or more interleaved repetitions could show shorter persistence. Further, we used large stimuli which induced strong gamma-band responses. Given that stimuli that generate stronger gamma-band responses initially also show stronger increases (see previous chapter), it is conceivable that the persistence of the effects also depend on the strength of gamma-band responses. The persistence of an adaptation effect does not necessarily show a positive relationship to stimulation intensity, however (Ganmor et al., 2010).

Many previous studies have used small stimuli, targeted to isolated single cells (Giaschi et al., 1993; Kohn and Movshon, 2003; Albrecht et al., 1984; Priebe and Lisberger, 2002; van Wezel and Britten, 2002; Movshon and Lennis, 1979). Given the strong, qualitative difference in effects dependent on stimulus size, this limits the comparibility to the current design. Furthermore, a lot of studies of adaptation or repetition are performed under anesthesia, which could eliminate some repetition effects altogether.

Studies concerned with changes in coordination rather than rates are rare. Pairwise neuronal correlations can be affected in ways that are poorly understood (reviewed in Solomon and Kohn 2014). An adaptation experiment by Jia and Kohn (Jia et al., 2011) in anesthetized monkeys specifically investigated adaptation of V1 gamma-band responses. Here, a prolonged presentation of a large, moving grating followed by repeated "top-up" presentations of the adapter yielded response decreases in response to the repeated grating when the grating direction was preferred by the recorded V1 location, or no change when it was not preferred. Effects recovered fully within minutes and no cumulative repetition increases were reported. Major differences compared to the repetition paradigms employed here and in Brunet et al. (2014) are the initial prolonged, uninterrupted exposure and the anesthesia.

In one of the few studies using awake animals, Stoelzel et al. (2015) reported a long-lasting effect of minutes of visual stimulation (for up to an hour) in the LGN on rate responses, of passively observing rabbits. As discussed above, the stimulation paradigm was different compared to the present study, in particular the stimulus size. However, since tests on this timescale are rather rare in awake animals, it illustrates that visual stimulation in awake animals (and across species) can result in plasticity on this timescale. Furthermore, particular forms of repetitive, artificial stimulation can lead to long-term potentiation (LTP) or depression of synapses, and consequently increased or decreased rate outputs also on a similar timescale. The observed effects strongly depend on the timing used in the stimulation protocols (Bliss and Lomo, 1973; Malenka and Bear, 2004). Particular forms of repetition effects in mice, with increased VEP and rate responses in V1 , have been linked to LTP (Cooke and Bear, 2010; Cooke et al., 2015; Cooke and Bear, 2012; Frenkel et al., 2006). However, the effects in mice only become apparent across days, likely involving sleep.

The opposite direction of observed rate effects and the reset of the observed gamma-band response strength with rest (Brunet et al. 2014 and Figure 3.7) suggests that the present form of plasticity differs from this exposure-based increase in mice. In the current paradigm, effects of rest (with a black background, and animals frequently closing their eyes) could be differentiated from effects of a local break in stimulation. Only a "proper" rest period resulted in a reset of gamma-band response strength. Such rest in a darkened environment, either by design or through the natural behavior of eye closing, means that the retina and thereby the entire visual system receives less input, and could show various forms of disadaptation. It is therefore unclear whether the effect of rest observed here is purely cortical, purely in the input structures, or both.

The repetition effect in the present study was also robust to "interference". Rapid, repeated presentation of grating stimuli different from the repeated stimulus did not result in a reset to an initial gamma-band response for the repeated stimulus (Figure 3.7). This is interesting because it certain forms of repetition effects, such as repetition suppression in monkey IT or the fly olfactory system, can be easily undone with such interference (Vogels, 2016; Hattori et al., 2017). Repetition suppression in flies can last on the order of 20 min to an hour after just 15 odor presentations. The presentation of other stimuli is thought to "normalize" responses, through a well-described dopaminergic synaptic modulation (Hattori et al., 2017). The current "interference" paradigm only constitutes a first attempt at this design. In the future, it would be interesting to extend this paradigm to include interference with stimuli that are more dissimilar to the repeating stimulus, and prolong the duration of the interference further. Similarly, the persistency after several minutes of stimulation at a different location only provides a lower bound for the duration of this effect, and longer breaks require testing. In summary, the present experiment demonstrated persistence of the repetition effect given several minutes of stimulation at a different location, or a few minutes of intervening stimuli. This is indicative of some form of plasticity, since V1 normally integrates information only on the sub-second timescale (Kim et al., 2019; Ringach and Shapley, 2004). The current study used a massed repetition design, without intervening stimuli as in the previous chapter. This maximized the number of blocks obtainable for slope fitting, and allowed a relatively straightforward interpretation, because repetition effects could not have been influenced by another stimulus. It is possible that this form of massed, uninterrupted form of repetition is optimal to strengthen repetition effects, such that effects observed here may be unusually strong or able to generate strong memory effects. On the other hand, it is possible that a counteraction by a temporally local form of adaptation was particularly strong in this case, so that repetition effects and their persistence could be even stronger with a different paradigm.

3.6 Chapter summary

In this chapter, it was shown that with the massed repetition (50 times and more) of a stimulus, the gamma-band increase with repetition

1. is position specific and therefore likely not the result of top-down feedback from higher-order areas and
2. shows memory on the timescale of minutes, given continuous task engagement.

Chapter 4

Opposing effects of spatial predictability on gamma and firing

4.1 Acknowlegdements for Chapters 4 and 5

The following two chapters are based on Peter et al. (2019). The publication reflects a collaborative effort between the labs of Martin Vinck (Martin Vinck, Cem Uran), Pascal Fries (Pascal Fries, myself (Alina Peter), Rasmus Roese, Jarrod Dowdall), and Wolf Singer (Wolf Singer, Johanna Klon-Lipok, Sylvia van Stijn and William Barnes). I, Cem Uran and Martin Vinck conceived of the idea of the study and designed the experiments. I, Cem Uran, Johanna Klon-Lipok and Rasmus Roese performed recordings. I, Johanna Klon-Lipok, Rasmus Roese, Sylvia van Stijn and William Barnes performed initial behavioral training. Johanna Klon-Lipok, Sylvia van Stijn, William Barnes and Wolf Singer planned and performed surgical implants. For this I wish to acknowledge also Michael Schmid and Richard Saunders. Martin Vinck and Pascal Fries provided supervision.

I, Cem Uran and Martin Vinck performed data analysis. In particular, I performed all LFP- and MUA-LFP analyses and preliminary versions of the spiking analyses. Cem Uran produced final spiking analyses, developed the final algorithm for receptive field analyses, and developed the peak-fitting algorithm with the help from Martin Vinck and me. Martin Vinck performed the regression analyses for Figure 5.9. All figures presented here were produced by me, with input from Cem Uran and Martin Vinck for the indicated analyses.

All introduction and discussion sections in these chapters were written by me. The introduction and discussion in Chapter 5 is more elaborate compared to the original paper, which stressed the spatial prediction component. The methods are taken over *verbatim* from the paper, since these require precision rather than interpretation or contextualization. Similarly, the results are largely taken over *verbatim*, with the exception that supplemental figures were integrated into the main text, which required some small changes in the order of paragraphs and naming of figures. I will reiterate this fact at each respective section, but I will avoid self-citations in-text as would be common for smaller or less frequent text sections for the sake of readability.

4.2 Introduction

In Chapter 2 (Stimulus specificity of repetition effects), natural images induced highly stimulus-dependent levels of both gamma-band activity and firing rates. Gamma-band amplitude for the large, colorful natural images employed in Chapter 2 was prominently related to hue in all animals, with red colors inducing particularly strong responses (Figure 2.6C). Additionally, repetition effects were stronger for stimuli that induced relatively strong gamma (Figure 2.12). To better understand these gamma repetition effects and gamma in general, the underlying stimulus aspects determining gamma response strength need to be identified.

One central determinant of gamma appears to be stimulus size: gamma-band activity for achromatic grating stimuli is jump-started as soon as the stimulus starts to cover not just the classical RF of a particular neuron, but also its surround (Gieselmann and Thiele, 2008). Similarly, long achromatic bars that stimulate both RF center and surround induce gamma in V1 LFPs, whereas short ones do not (Chalk et al., 2010). Compared to the dependence of gamma on other stimulus features, this dependence on the surround may be a qualitative, instead of quantitative change.

Based on these findings, it has been theorized that gamma crucially depends on contextual interactions that yield surround suppression and sparse coding (Jadi and Sejnowski, 2014). Given the relationship between surround suppression magnitude for rates and the homogeneity between the stimulus falling on the classical RF and the surround (Coen-Cagli et al., 2015), it has been further suggested that gamma may crucially depend on predictable center-surround relationships (Vinck and Bosman, 2016). Besides stimulus size, gamma strength also depends on various other stimulus features such as contrast, orientation, motion speed and spatial frequency for well-controlled, achromatic stimuli such as gratings and bars (e.g. Hadjipapas et al. 2015; Henrie and Shapley 2005; Ray and Maunsell 2010; Frien et al. 2000; Jia et al. 2011; Orekhova et al. 2015; Muthukumaraswamy and Singh 2013). This means that the observed variance of gamma for different natural images is not entirely unexpected. For natural images and during natural vision, both presence and strength of gamma-band activity is a matter of an ongoing debate (e.g. Hermes et al. 2015; Brunet et al. 2015; Hermes et al. 2019). Chromatic (i.e. colored) compared to achromatic (black-and-white) natural images in the study of Brunet et al. (2015) seemed to yield stronger responses, although no systematic comparison was made in this study. Another study employing only achromatic natural images reported high variance in induced gamma-band amplitudes between images, and weak or absent gamma-band responses for many of the images (Hermes et al. 2015, but see Brunet and Fries 2019). In contrast, Rols et al. (2001) and Shirhatti and Ray (2018) reported prominent gammaband oscillations in the LFP for uniform chromatic stimuli. Together with the present natural image data (Chapter 2), this indicates an important role of color for gamma. This chapter studies determinants of gamma-band activity using uniform surface stimuli as a vehicle. Chromatic and achromatic stimuli are contrasted, and the role of contextual surround interactions for the interplay between gamma response strength and firing rates is studied. The immediately following Chapter 5 will then investigate the dependence of gamma-band amplitude on particular hues and develop an understanding of hue-dependency based on hue-dependent adaptation mechanisms.

4.2.1 Relationship between gamma, surround modulation and sparse coding

Vision requires the integration of information over space, and insight into relationships between information in space. In part, this integration is purely feedforward, because neurons with small RFs converge onto neurons with larger RFs along the visual hierarchy (Felleman and Van Essen, 1991; Serre et al., 2005; Lamme and Roelfsema, 2000; DiCarlo et al., 2012). However, even in V1, neuronal responses to sensory inputs depend strongly on their spatio-temporal context (see also General Introduction, section 1.4.2). In many visual areas and V1 in particular, firing rates to a stimulus in a neuron's classical RF (CRF) can be increased or decreased by stimuli presented in their surround. Its surround is of approximately circular shape surrounding the CRF, and does not elicit responses when stimulated in isolation, i.e. without also stimulating the CRF (Spillmann et al., 2015; Vinje and Gallant, 2000; Rao and Ballard, 1999; Angelucci et al., 2017; Gilbert, 1992). Surround modulation is the result of lateral and feedback network interactions, which may inform a given V1 neuron about a larger image region than covered by its CRF (Lund et al., 2003; Gilbert, 1992; Angelucci et al., 2017). A rate decrease with surround stimulation is typically referred to as "surround suppression" and occurs for a wide range of stimuli. Numerous related computations have been linked to surround modulation, among them normalization (Carandini and Heeger, 2011), contour integration (Liang et al., 2017), perceptual filling-in (Zweig et al., 2015; Land, 1959; Wachtler et al., 2003), figure-ground separation (Lamme, 1995), and the computation of a saliency map (Coen-Cagli et al., 2012; Li, 2002). Importantly, a lot of these computations and their associated signature neuronal phenomena can be captured using the overarching goal of sparse coding. Sparse coding refers to the generation of informative neuronal responses that are as rare or low as possible in terms of the percentage of the neuronal population involved or the overall rates produced over time (Zhu and Rozell, 2013).

A sparse code is elementary to both efficient and predictive coding operations (Rao and Ballard, 1999; Vinje and Gallant, 2000). Theories of efficient coding propose that surround modulation removes image redundancies (e.g. homogeneous parts) across space from neuronal representations (Schwartz and Simoncelli, 2001; Simoncelli and Olshausen, 2001; Coen-Cagli et al., 2012, 2015; Rao and Ballard, 1999; Barlow, 2001; Vinje and Gallant, 2000; Zhu and Rozell, 2013). Theories of predictive coding in turn interpret the resulting neuronal response as a prediction error signal (Friston, 2005; Rao and Ballard, 1999; Spratling, 2010). Feedback and lateral connectivity, and therefore coordination, is required in any of these interpretations of surround modulation. Consequently, surround modulation may also affect temporal correlations among neuronal responses (Singer and Gray, 1995).

Whereas many theoretical frameworks typically focus on explaining firing rate modulations, recent work on efficient and predictive coding also includes considerations on the modulation of neuronal synchronization (Bastos et al., 2012; Jadi and Sejnowski, 2014; Chalk et al., 2016; Vinck and Bosman, 2016). Neuronal synchronization is relevant to the encoding and transmission of information (e.g. Buzsáki and Wang 2012; Vinck et al. 2010a; Havenith et al. 2011; Salinas and Sejnowski 2001; Sejnowski and Paulsen 2006; Fries 2005; Fries et al. 2007; Fries 2009; Börgers and Kopell 2008; Varela et al. 2001; Wang 2010; Buzsáki 2006; Singer 1999; Singer and Gray 1995; Bressler et al. 1993; Bernander et al. 1994; Abeles 1982; Azouz and Gray 2000; Akam and Kullmann 2014). Therefore, it could play an important role in contextual integration. V1 gamma-band synchronization in particular appears linked to surround

modulation because of the repeated finding that the amplitude of V1 gamma oscillations increases with stimulus size (Gieselmann and Thiele, 2008; Ray and Maunsell, 2011; Jia et al., 2013b; Chalk et al., 2010; Perry et al., 2013; Jia et al., 2011; Gray et al., 1990). For a detailed discussion, refer to Vinck and Bosman 2016. Perspectives differ on the relationship between gamma oscillations and predictive or efficient coding operations and, consequently, center-surround relationships (Jadi and Sejnowski, 2014; Bastos et al., 2012; Arnal and Giraud, 2012; Vinck and Bosman, 2016; Chalk et al., 2016; Korndörfer et al., 2017). Some modeling studies suggest neuronal synchronization as a hallmark of an efficient code (Jadi and Sejnowski, 2014; Chalk et al., 2016).

Bastos et al. (2012) and Arnal and Giraud (2012) hypothesized that whereas lower frequency bands carry feedback predictions from higher areas, gamma-band activity aides in the encoding and feedforward transmission of prediction error signals (Bastos et al., 2012, 2015; Arnal and Giraud, 2012). In accordance with this hypothesis, based on multi-area recordings along the visual hierarchy, it was found that bottom-up directional influences are strongest in the gamma, but top-down influences in the lower alpha/beta (≈10-20 Hz) band, respectively (Bastos et al., 2015; Richter et al., 2018; van Kerkoerle et al., 2014; Bressler et al., 2006; Bosman et al., 2012; Michalareas et al., 2016). From this hypothesis, it follows that a mismatch between center and surround stimuli should lead to an increase in both firing rates and gamma-band amplitudes, signaling prediction errors to higher areas.

Alternatively, Vinck and Bosman (2016) recently hypothesized that the amplitude of gamma oscillations reflects the degree to which classical RF inputs are predictable from the surround. Gamma synchronization between mutually predictable neuronal populations could provide a mechanism for orchestrating interactions between distributed neuronal populations (see Discussion). According to this "predictability hypothesis", a mismatch between center and surround stimuli should result in an increase of firing rates but a decrease of gamma-band responses. Of note, the two hypotheses cannot be distinguished based on the rate responses. To distinguish between these conflicting views, center-surround predictability was manipulated in the following experiments. As a starting point to test the relationship between center-surround predictability, gamma synchronization and rate responses, uniform surfaces were used. Uniform surfaces contain identical, and therefore redundant, information across a relatively large part of the image.

4.2.2 Processing of chromatic versus achromatic stimuli

Uniform surfaces are defined by their hue, luminance and saturation. Surfaces can be chromatic or achromatic (black and white). In either case, such surfaces have a high degree of predictability at the physical image level. Nevertheless, there are likely substantial differences in the way these surfaces are processed by area V1. As a consequence, their "predictability" on the neuronal level may also differ strongly. In particular, there are two different ways through which V1 surface responses may arise (Zweig et al., 2015).

1. Local: Neurons with RFs at the uniform center of a surface may be driven by direct feedforward inputs. These may generate redundant (predictable) signals locally. Specifically for chromatic surfaces, a particular type of hue-selective neuron with RFs at the uniform surface region may directly encode color and luminance information. These are so-called single-opponent neurons (in LGN

and/or V1). They possess L+/M-, M+/L-, or blue (S) and yellow (L and M) color opponencies that respond to *chromatic* surfaces rather than edges (Shapley and Hawken 2011; Livingstone and Hubel 1984, for an introduction to color, see General Introduction, section 1.4.1, Color in V1). Moreover, Chen et al. (2007) reported that there is an additional, strong source of drive for chromatic compared to achromatic brief flashes of homogeneous light.

2. Edge-derived: Surface information could be initially derived by neurons with RFs at the edge of the surface, and via lateral interactions inform neurons with RFs more toward the center (Zweig et al., 2015; Land, 1959; Wachtler et al., 2003).

The relative contributions of these two pathways (local vs. edge-derived) remain largely unknown. Importantly, they likely differ between chromatic and achromatic surfaces (Zweig et al., 2015; Zurawel et al., 2014). In particular, Zweig et al. (2015) compared responses at the surface's center to the edge. Responses to achromatic surfaces were consistent with an edge-derived "fill-in" mechanism. In contrast, this was not observed for chromatic surfaces (Zweig et al., 2015). Possibly, this was precluded by the availability of surface-based information provided by single-opponent cells. This indicates that contextual interactions likely have a different nature for achromatic than chromatic stimuli. The following section will test for differences in contextual modulation of firing activity and gamma-band synchronization between chromatic and achromatic surfaces, using stimuli of different sizes and a center-surround mismatch paradigm.

4.3 Methods

Preliminary remarks: Since the description of methods requires precision and leaves little room for variations in expression, the methods section is taken over ***verbatim*** *from the respective publication, Peter et al. (2019). See Chapter 4, section 4.1 for acknowledgement of methods development. Note that monkey A in this study is not the same animal as monkey A in the previous chapters (but comes from the same lab), whereas monkey H is the same animal.*

All procedures complied with the German and European regulations for the protection of animals and were approved by the regional authority (Regierungspräsidium Darmstadt).

4.3.1 Surgical procedures

Two male adult macaque monkeys (Macaca mulatta) were used in this study (age 9-10 years, 15-17 kg). All surgeries for implantations were performed under general anesthesia and were followed by analgesic treatment post-operatively. A head post was implanted in both monkeys to allow for head fixation. In monkey H, we implanted CerePort ("Utah") arrays with 64 microelectrodes (inter-electrode distance 400 μm, tip radius 3-5 μm, impedances 70-800 kOhm at 1000 kHz, half of them with a length of 1 mm and half with a length of 0.6 mm, Blackrock Microsystems). One such array was implanted into area V1, another one in V4, both in the left hemisphere. The V4 array is not considered here. For array implantation, a large trepanation covering both areas was performed, the dura was cut open and reflected, arrays were inserted using

a pneumatic device (Blackrock Microsystems), and both dura and bone were surgically closed. A reference wire was inserted under the dura towards parietal cortex. In monkey A, we implanted a semi-chronic microelectrode array Microdrive into area V1 of the left hemisphere (SC32-1, Gray Matter Research, containing 32 independently movable Alpha Omega glass insulated Tungsten electrodes with an impedance range of 0.5-2 MegaOhm and an inter-electrode distance of 1.5 mm). The microdrive chamber was used as the reference during recordings. The precise layers/depths that were recorded from could not be identified based on histological verification, which is the current gold-standard, because the animals are still alive. However, based on the observation that all sites in monkey H and the vast majority of sites in monkey A do not show the typical inversion of the event-related potential as is found in the deep layers (Li et al., 2015), we estimate that our recordings mainly sample activity from layers 2-4. Sites in monkey A and monkey H behaved qualitatively in a consistent manner across depths, such that all recording sites were pooled.

4.3.2 Behavioral task

Both monkeys were trained on a fixation task. Monkeys were seated in a custom-made primate chair in a darkened booth. The two animals were positioned 83 (monkey H) or 64 cm (monkey A) in front of a 22 inch 120 Hz LCD monitor (Samsung 2233RZ, Ghodrati et al. 2015; Wang 2011). Both monkeys self-initiated trials by fixating on a small fixation spot, which was presented at the screen center. Monkey H performed a pure fixation task. For monkey H, the fixation spot was a Gaussian with a white center, tapering smoothly into the background. For recordings with white background, the fixation spot color was changed to red. Note that the pattern of results for gray and white backgrounds was very similar despite this difference (Chapter 5, Figure 5.8), and that receptive fields were not covering the fovea. The task of monkey A was to report a change in the fixation spot from red to green or blue (randomly) with a lever release. The change in the fixation spot occurred only after the stimulus period and an additional 700 ms of background stimulation, during which the animal maintained fixation. For the recordings with colored backgrounds in monkey A, fixation colors were changed to remain visible, with a magenta fixation spot during the baseline and stimulus period. For both animals, trials during which the eye position deviated from the fixation spot by more than 0.8-1.5 dva radius were aborted. Correctly performed trials were rewarded with diluted fruit juice delivered with a solenoid valve system.

4.3.3 Recordings

Data acquisition was performed using Tucker Davis Technologies (TDT) systems. Data were filtered between 0.35 and 7500 Hz (3 dB filter cutoffs) and digitized at 24.4140625 kHz (TDT PZ2 preamplifier). Stimulus onsets were recorded with a custom-made photodiode. Eye movements and pupil size were recorded at 1000 Hz using an Eyelink 1000 system (Eyelink Inc.) with infrared illumination. Eye signals were calibrated before each recording session using a standardized fixation task. Behavioral control and stimulus presentation was done using in-house custom software running in Matlab, including ARCADE (Dowdall et al., 2018).

4.3.4 Visual stimulation paradigms during recordings

For all paradigms, stimuli were circular, did not have overlap with the fixation spot, and typically spanned a region from ca. 3-9 dva of eccentricity (monkey H) or 2.5-8.5 dva (monkey A, maximum: 1.6-9.6 dva for *Dataset 4*) in the lower right visual quadrant, matching RF locations. Trials always started with a baseline that lasted 0.5-0.6 s (monkey H) or 0.5-0.8 s (monkey A), and during which only the FSB and the fixation spot was shown. We used the following stimulus paradigms:

Dataset 1: For Figures 4.2, 5.1B and 5.6, we presented large uniform stimuli of 6 dva diameter on a gray FSB. For the chromatic condition, we used stimuli that were either green, red, or blue, at three different luminance levels (which are shown in Figure 5.1). For Figure 4.2, only the chromatic conditions with the highest available luminance level were used, approximately corresponding to the maximum possible luminance level for the blue primary. For the achromatic condition, we used either black (minimum luminance) or white (maximum luminance) stimuli.

The background was of an intermediate gray value that allowed for good eye tracking quality (see Supplementary Table 5.5 at the end of Chapter 5 for all luminance and CIE values). Stimulus duration was 3.3 s. This dataset included 3 sessions from monkey H and 2 sessions from monkey A. There were 20±0 (H) and 20±0 (A) trials in each session for each of the 11 conditions (2 color hues * 3 luminance levels + black and white).

Dataset 2: For Figure 4.5, i.e. the size tuning paradigm, we presented a smaller (either 0.5, 1, or 2 dva) stimulus and a larger (6 dva) surface stimulus in the same trial sequentially, with each stimulus presented for only 0.6 s. In each trial, either the smaller ("small-first") or largest ("large-first") surface was presented first. In addition, we used an "edge" condition in which the selected multi-unit's RF was centered around the vertical edge of the 6 dva stimulus, again followed or preceded by the standard full condition (Figure 4.8). The colors used were red, blue and green (at the same luminance intensities shown in Figure 4.2), black and white, and in case of monkey H, also orange, cyan and magenta hues. This dataset included 5 sessions from monkey H and 4 sessions from monkey A. There were 12.78±4.3 (H, 64 conditions) and 12.86±5.4 (A, 40 conditions) trials in each session for each of the conditions (4 stimulus sizes * 2 presentation orders * 8/5 colors (H/A)).

Dataset 3: For Figure 4.7, we used only red, green and blue hues (with the same luminances as the maximum luminant red, green and blue used in *Dataset 1*, Figure 4.2). We presented three stimulus conditions: The uniform surface, the "annulus" and the "blob" condition (Figure 4.7). Stimuli in annulus or blob conditions were of the same size as the uniform surface, but the center 1 dva of the surface was either surrounded by a thin (0.25 dva) annulus of one of the other, equiluminant, hues, or filled completely with one of the other hues (Figure 4.7). For each surface of a given hue, there were therefore two "annulus" and "blob" conditions with the two remaining colors (Figure 4.7). In the analysis, we averaged over all the color combinations for a given condition, and compared the three main conditions. For monkey H, we additionally recorded two sessions with maximally luminant instead of equiluminant hues. Note that this generated strong luminance contrast changes between the colors, but yielded qualitatively similar results. This indicates that the observed effects do not depend on equiluminance, a condition that may occur rarely in nature. Because results were qualitatively similar, we pooled these sessions with the remaining 5 sessions of this animal. We used stimulus presentation times of 1.3-3.3 s. The first 1.3 s were analyzed, as in Figure 4.2. This dataset included 7 sessions from monkey H and

1 session from monkey A. There were 15.88±0.21 (H) and 18.87±0.34 (A) trials in each session for each of the 15 conditions (3 color hues * 2 color hues for mismatch * 2 (annulus vs blob mismatch) + 3 uniform conditions).

Dataset 4: For Figure 5.1A, we recorded "rainbow" sessions in which surfaces (again 6 dva diameter size) of different colors were presented at the maximum possible luminance. We sampled the visible light spectrum linearly in 15 steps of equal size in terms of wavelength, with the MATLAB (MathWorks, Inc.) internal function spectrumRGB.m. Note that the monitor cannot produce line spectra, but can only approximate the corresponding hues through mixing of RGB channels (see e.g. Figure 5.9A for a yellow hue). We additionally included brown and pink (extra-spectral) hues and achromatic stimuli (see Table 5.5 and 5.2). This dataset included 3 sessions from monkey H and 2 sessions from monkey A. There were 12.1±6.9 (H) and 20±0 (A) trials in each session for each of the 22 conditions. For the analyses shown in Figure 5.1B, we used *Dataset 1*.

Dataset 5: For Figures 5.7 and 5.9, we used FSBs of various hues. The backgrounds used were red, green, blue and yellow at maximum possible luminance, as well as black, white and gray, presented at the same luminance intensities as in the other datasets. Surface stimuli of 6 (monkey H) or 8 (monkey A) dva diameter in size were used. The size was slightly increased for monkey A to place the edge of the surface stimulus further from the most peripheral RFs. The hues used for the surface were identical to the ones used for the FSBs. In addition, we presented chromatic surfaces with reduced values, namely red, green and blue with the same luminance levels as in Figure 4.2, and a brown surface. All possible combinations of surface and FSB hues were shown. All other presentation parameters were kept as for *Dataset 1*.

For all stimulus paradigms for monkey A, and in *Dataset 5* for monkey H, there was a post-stimulus period of 0.7 s (0.5 s in monkey H) after the offset of the stimulus, during which the monkey was required to maintain fixation. For monkey A, the fixation color would change after this period and the monkey Had to respond to this change with the release of a lever, whereupon the fixation spot was removed. Presentation of different stimulus conditions was in a pseudo-random order. This dataset included 16 sessions from monkey H and 9 sessions from monkey A (1-2 per FSB). There were 17.89±0.17 (H) and 19.00±0.08 (A) trials in each session for each of the 11 conditions.

4.3.5 DKL Color Space

In order to calibrate the monitor outputs, the luminance of the RGB monitor primaries were measured with Konica Minolta CS-100A chroma meter and look-up tables were generated. Monitors were gamma-corrected to linearize the dependence of luminance on RGB values. The Derrington-Krauskopf-Lennie (DKL) Color Space was introduced as a color-opponent modulation space (Krauskopf et al., 1982; Derrington et al., 1984). DKL color space is based on a cone-contrast representation, where cone activation to a color stimulus is quantified as the relative change of the cone activations with respect to the background color (Brainard, 1996). Weber cone-contrasts are computed in three steps: 1) The change in cone-activation relative to the full-screen background is computed, 2) This change in cone-activation is normalized (divisively) by the extent to which the background differentially activates the different cones. These cone contrasts are then transformed into 3 primary axes of the DKL space, which correspond to the mechanisms of L+M (luminance), L-M (red-green opponency), and S-(L+M) (blue-yellow opponency) as described in Brainard (1996). Along the L-M axis, max-

imum L/M cone contrasts were 9.60% and 14.81% respectively, along the S-(L+M) axis, S cone contrast was 79.35%. These values were found to be similar to previous studies (Hansen and Gegenfurtner, 2013; De Valois et al., 2000).

4.3.6 Data analysis

Preprocessing. Data were analysed in Matlab using the FieldTrip toolbox (Oostenveld et al., 2011). Only correctly performed trials were analyzed. LFPs were derived from the broadband signal using Matlab's decimate.m function, by low-pass filtering with a cutoff frequency of 24414.0625/24/2 Hz (FIR Filter with order 30) and downsampling to 24414.0625/24 Hz. Line noise was removed using two-pass 4th order Butterworth bandstop filters between 49.9-50.1, 99.7-100.3 and 149.5-150.5 Hz. LFPs had a unipolar reference scheme described in *Recordings*. Explorative analyses with local bipolar derivations, obtained by subtracting the signals from immediately neighboring electrodes from each other, yielded comparable results (data not shown). MU (multi-unit) signals were derived from the broadband signal through bandpass filtering between 300 and 6000 Hz (4th order butterworth), rectification, and applying low-pass filtering and downsampling the same way as for the LFPs. For the calculation of rate modulations, this MU signal was smoothed with a Gaussian kernel with an SD of 20 ms. This signal is called MUA in previous chapters, but will be typically referred to as MU signal or MU firing in these chapters, in keeping with the original publication and the figures. This was deemed preferable to making small changes to a number of figures, and cluttering the text with pointers to this. Qualitatively similar results were obtained using thresholded multi-unit data. We used this MU signal for all analyses in the main text, as in previous studies by other labs (Schmid et al., 2013; Self et al., 2013; Xing et al., 2012; Legatt et al., 1980).

Receptive field estimation. Receptive fields were mapped with moving bar stimuli (spanning the entire monitor). Moving bars (width 0.1 dva, speed 10/17 dva/s) were presented in 8 orientations for monkey H and 8-16 orientations for monkey A, each for 10-20 repetitions. Mapping sessions were intermittent for monkey H and typically daily for monkey A, to confirm stability of the recordings. MU responses were projected onto the stimulus screen, after shift-correction by the response latency that maximized the back-projected response. MU responses were then fitted by a Gaussian function. This Gaussian was used to extract the 10^{th} percentile and the 90^{th} percentile, and this was done separately for each movement direction. Across the 16 directions, this yielded 32 data points, which were fit with an ellipse. This ellipse was defined as that MU's RF. The RF size is defined as the diameter based on (area of the ellipse/pi)*2.

Electrode selection. We included all electrodes for analysis that met the following criteria: (1) the MU showed a response to RF stimulation that was at least two SDs above stimulation outside the RF. (2) The MU response during the response period (0.05-0.15 s) of at least one condition of the respective dataset was at least 2 SD above the corresponding baseline (-0.1-0 s). In case of Figures 4.5-4.7, it was additionally required that the RF center of the MU was within 0.5 dva of the stimulus center. In the remaining figures, it was required that the RF center was within the surface stimulus.

Estimation of LFP power spectra. For Figures 4.2, 4.7-5.1 and 5.7-5.9, the baseline period was the last 500 ms before stimulus onset, and each stimulation period yielded two non-overlapping epochs of 500 ms (0.3-1.3 s period). For Figure 4.5, due to the short presentation times, we used epochs of 300 ms (300-600 ms after the onset of the stimulus, and for baseline 300 ms before stimulus onset). LFP epochs were multi-

plied with discrete prolate spheroidal sequences (multi-tapers for ±5 Hz smoothing), Fourier transformed and squared to obtain LFP power spectral densities (for a recent discussion on spectral estimation see Pesaran et al. 2018). For Figure 5.6, we used windows of 0.3 s length, slid over the data in steps of 50 ms. Data were multiplied with a Hann taper before Fourier transformation.

Normalization of LFP power spectra. To show LFP power changes, we computed relative power spectra by dividing single-trial power spectra from the stimulation period by the average power spectra across conditions and trials from the baseline. This was shown as a fold-change in all figures showing relative changes except for Figure 5.6, the time-frequency representations (TFRs), where for visualization purposes, we transformed this into dB units.

To investigate absolute LFP power (without reference to the baseline), we normalized power spectra per electrode by the total power above 25 Hz in the baseline condition. This normalization reduced variance or scaling in the LFP power spectra across sessions and animals before averaging. By normalizing both the baseline and the stimulus period by the same normalization factor, we could still examine changes in raw LFP power across conditions, for each frequency bin separately. This would not have been possible if we had normalized the LFP power spectrum in a given condition by the total power across frequencies in the same condition. These power spectra were averaged across the selected channels (except for single-channel analyses as in Figure 4.4).

Quantification of LFP gamma-band amplitude. Quantification of the differences in gamma-band amplitude between conditions is in general a difficult problem because changes in firing rate can cause broad-band shifts in the LFP power spectrum, and because spikes can "bleed-in" at higher LFP frequencies (Miller et al., 2009a; Ray and Maunsell, 2011; Pesaran et al., 2018; Buzsáki et al., 2012). We developed an algorithm to extract gamma-band amplitude in order to address these problems (see Figure 4.1 for an illustration). We present two versions of this algorithm that are

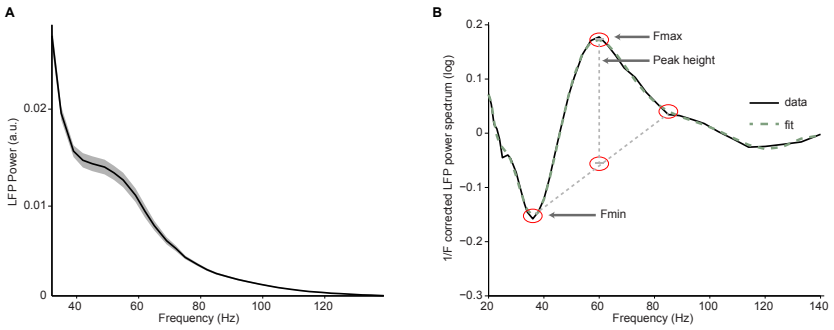

Figure 4.1 — Illustration of fitting procedure. (A) Average LFP power spectra for a large chromatic condition of an example session used in Figure 4.5. LFP spectra for all conditions were normalized to the summed power (> 25 Hz) for the baseline condition. (B) Log-transformed, $1/F^n$ corrected spectra (solid line) and their fit (dashed line). Peak height was determined as the difference between the peak value at location Fmax and a baseline estimate based on the average of the power at location Fmin and Fmin+2*(Fmax-Fmin), the estimate of peak width.

used for separate figures, and are based on constructing a polynomial fit of the LFP spectrum which was detrended in two separate ways. The first algorithm had the following structure:

1. Power spectra were log-transformed and the frequency axis was also sampled in log-spaced units to avoid over-fitting of high-frequency datapoints. All subsequent polynomial fits were performed on the 20-140 Hz range.

2. We used the change in stimulus-induced LFP power versus the common baseline (see above), expressed as $\Delta P = \log(P_{stim}) - \log(P_{base})$.

3. To determine the polynomial order, we used a cross-validation procedure to prevent overfitting. A random half of the trials was used for the fitting and deemed the "training set". The remaining trials were the "test set". Polynomials of order 1-20 were fit to ΔP as a function of frequency for the "training set", minimizing the mean squared error. We then computed the mean squared error using the same polynomial fit on the "test set" for each of the 20 orders. This procedure was then repeated for multiple (50) iterations, with a random half of the trials selected for each iteration, and for each iteration, the best-performing order was retained.

4. A polynomial with the median of the best-performing orders was then fit to the complete set of trials.

5. On the polynomial fit, local maxima and minima in the 30-80 Hz range were identified. The peak gamma frequency was the location of the maximum. The band-width of gamma was estimated as twice the distance between the frequency of the maximum (F_{max}) and the frequency of the first local minimum to the left of the maximum (F_{min}), i.e. $b = 2F_{max} - F_{min}$ (see Figure 4.1). The gamma amplitude was then assessed from the difference between the value of the polynomial fit at the maximum and the average of the polynomial fit at F_{min} and $F_{max} + F_{min}$ (Figure 4.1).

6. This difference was taken in log-space (because the power spectra were originally log-transformed) and then transformed to a fold-change.

As in previous chapters, log refers to a base of 10. If firing rate changes relative to baseline (or between conditions) were very strong, e.g. with small stimuli, this fitting procedure occasionally ran into problems, because relative LFP power spectra showed broad increases that were likely due to non-rhythmic processes like spikes or postsynaptic potentials (see Figure 4.7C for an example of this effect). In addition, in Figure 5.7 and 5.9, because we used background stimuli of different hues, a "neutral baseline" like the gray background screen was not always available. In these cases we modified the second step of this algorithm. Instead of computing the change in LFP power relative to baseline, we performed a $1/F^n$ correction on the raw LFP power spectrum. The $1/F^n$ correction was performed by fitting an exponential to the LFP power spectrum, excluding data points in the typical gamma range of 30-80 Hz. Note that we fitted an exponential function because in many cases, bleed-in of spiking energy in the LFP caused a departure from a linearity in the log(power) versus log(frequency) graph (see also Haller et al. 2018; Shirhatti and Ray 2018). We visually inspected the fits for a large number of spectra and compared this also to a procedure with a mixture of a linear fit and a Gaussian fit to the log(power) versus log(frequency) graph, which

had substantially more problems in dealing with spike-bleed at high frequencies, as well as with additional peaks (potentially harmonics) at higher frequencies (e.g. for the red surfaces, data not shown).

Spike-field coherence. For spike-field coherence, we used only electrodes selected by the procedure described above. In addition, for MU-LFP pairs, we required that the electrodes were direct neighbors in the grid, and in the case of monkey H, given that the microelectrode array had two fixed depths, were of the same depth. Spike-field phase-locking was computed as follows. We estimated the cross-spectral density between LFP and MU signal for each trial separately (cross-spectra) using the same spectral estimation settings as for the LFP power spectrum. This yielded one cross-spectrum per trial. We then normalized the cross-spectrum per trial by its absolute values, to obtain the cross-spectral phases (without amplitude information). We used those normalized cross-spectra to compute the Pairwise Phase Consistency (PPC), using FieldTrip (Oostenveld et al., 2011). This measure has the advantage that the bias by trial count, inherent to e.g. the spectral coherence, is avoided (Vinck et al., 2010b). For a given MU site, the PPC values were then averaged across all the combinations with LFPs from the other selected channels. Note that MU-LFP combinations from the same electrode were excluded to avoid artifactual coherence due to bleed-in of spikes into the LFP (Ray and Maunsell, 2011; Buzsáki et al., 2012). Because of the distance between electrodes (at least 400 μmeter), this was not an issue for MU-LFP combinations from different electrodes. The standard error of the PPC was estimated across sessions. This was different from SE estimation for power and rate, which used the bootstrap (see below). Bootstrap estimates are problematic for PPC because bootstraps contain repetitions of identical trials, which trivially yield high coherence values.

Rate modulation. Rate modulation was computed as $\log(M_{stim}/M_{base})$, where M_{stim} and M_{base} represent the MU firing activity in the stimulus and baseline period, respectively. To quantify surround suppression, we took the differences of these rate modulation indices between small and large stimulus size conditions.

Modulation index of fold-changes. To quantify the modulation of LFP gamma-amplitude (expressed as fold-change) between conditions (Figure 5.6, 5.7), we computed a modulation index as $(A-B)/(A+B)$, where A and B are the gamma-amplitudes in the two conditions, taken as the fold-change minus 1. Note that the fold-change was extracted using the polynomial fitting procedure described above, and a fold-change of 1 indicated the absence of a gamma peak.

Microsaccade detection and subsequent LFP analysis. For microsaccade detection, we smoothed horizontal and vertical eye signals (rectangular window of ±5 ms) and differentiated the signals over time points separated by 10 ms to obtain robust eye velocity signals. For monkey H, for whom data from both eyes were available, data were averaged across eyes. We then used the microsaccade detection algorithm described in Engbert and Kliegl (2003) with a velocity threshold of $6*c$, where c is the criterion defined as $c = \text{Median}[v^2] - (\text{Median}[v])^2$. Threshold crossings in either the horizontal or vertical direction were considered as microsaccades. We tested several threshold levels and obtained qualitatively similar results. We then removed data epochs of 100 ms after each microsaccade and recomputed our analyses (based on Lowet et al. 2016; 100 ms is approximately the duration of microsaccade effects in V1). Removing 200 ms after each microsaccade yielded qualitatively similar results but fewer remaining epochs. For the analysis of LFP gamma amplitude, we switched to analyzing epochs of 100 ms using a Hann taper, instead of the 500 ms time bins used before. This is sacrificing some frequency resolution and limiting the results to

frequencies > 20 Hz, in order to obtain a large number of microsaccade-free epochs. Epochs were zero-padded to 1 s, effectively smoothing the spectra. Note that we show the results for the data including microsaccades with the identical epoch length and taper to allow a fair comparison.

Pupil responses. Pupil signals across the two eyes were averaged for monkey H. Pupil size during the comparatively stable period 200 ms to stimulus onset was used as a baseline. Pupil size was then computed as percent change from the average response during this time (A-B)/B, where A is the pupil response at each time point and B is the average response during the baseline period. Note that since the Eyelink system gives outputs with arbitrary units, and these were negative during the baseline period, we took the absolute value for the denominator such that pupil size decreases are indicated by negative values.

4.3.7 Statistics

Error bars or shaded error regions correspond to ± one SEM. SEM was estimated using a bootstrap procedure, with the exception of spike-field coherence (see above). For the b-th bootstrap out of $B = 1000$ bootstraps, $b = 1, \ldots, B$, the following was done. For each condition in a given session, with a set of N trials $\mathcal{T}$, we took a random set of N trials from $\mathcal{T}$ with replacement. For that sample of N trials $\mathcal{S}_b$, we then computed the statistic of interest. For LFP signals, we then computed the average statistic in a given session over all channels, then averaged over sessions, and then monkeys. The rationale behind averaging across all LFP channels was that these signals are likely highly statistically dependent because of volume conduction among the relatively closely spaced electrodes. For MU signals, we computed the average statistic of interest across sessions per MU site separately, and then averaged across all recording sites. The standard error of the mean was then defined as the standard deviation over the B average statistics, as is common with bootstrapping procedures.

We used the bootstrap distributions for inference on fold-change estimates or fold-change modulation indices between conditions, as well as differences in peak gamma frequency. In this case, we computed for each bootstrap the difference between average statistics for two conditions, and then tested whether this distribution was different from zero (with Bonferroni correction for number of comparisons).

For frequency- or time-resolved differences (in absolute and relative LFP power spectra and rate modulation scores), we used multiple-comparison corrected permutation tests: In this case, we shuffled the trials between two conditions per permutation P times, and then constructed a permutation distribution of average absolute differences between conditions. We equalized trial numbers for each comparison, for example between chromatic/achromatic conditions or the different stimulus sizes. We then compared the observed difference between average statistics against this permutation distribution. For multiple-comparison correction, we used the procedure from Korn et al. (2004), which is based on the sorted distribution of absolute differences, with alpha and false discovery rate values of 0.05. In this iterative procedure, values in the observed distribution exceeding the 95th percentile of the P maximal values of each permutation distribution (critical value) are deemed significant. Significant values are removed from the observed distribution, and the same positions are removed from all P permutation distributions. Values in the observed distribution exceeding the critical value based on these permutation distributions are then iteratively collected until no value in the observed distribution exceeds the critical value. Note that statistical parameters are reported mostly in the figure captions.

4.3.8 Quantitative model for dependence of gamma-band amplitude on background stimulus

Cone data (bleaching difference corrected spectra) were extracted from Hárosi (1987). Polynomials of order 7 were fit to these curves. The cone response curves were then normalized to the maximum. We measured the spectral energy of each color as well as black, white and gray (Ocean Optics WaveGo; XWAVE-STS-VIS-RAD). The spectral energies of the colors were normalized to unit mass. For gray, we added the normalized energies of R, G and B and multiplied with the energy ratio of gray over white. We then convolved the cone response curves with the normalized spectral energies to determine how strongly each background adapts the three cones. Regression models were then fit as explained in the Results text and caption of Figure 5.9. SEM for regression coefficients were obtained by the same bootstrap procedure as described above.

4.4 Results

Preliminary remarks: This result section is based in large part, directly and verbatim on the results section of the respective paper, Peter et al. (2019). Some exceptions exist because the paper was subdivided into two chapters, i.e. 4 and 5. Furthermore, supplementary figures from the paper were either integrated into the text or omitted, with respective adjustments to the text. Figure numbers do not match those of the publication. It will become apparent to the reader that the experiments follow a logical order, building on each other, such that there is little point in changing the overall order. Similar to the methods section, there is only very limited choice in wording for large parts of these results, with little intellectual achievement if this were done. I therefore decided to quote the results section in its original form.

We recorded multi-unit (MU) activity and local field potentials (LFP) from the primary visual cortex (area V1) in two macaque monkeys, while they performed a fixation task. These recordings were made using a 64-channel chronic microelectrode array in monkey H and a 32-channel semichronic microelectrode array in monkey A (see Methods). Classical receptive fields (RFs, referring to classical RFs unless otherwise mentioned) of the MU activity were estimated using moving bar stimuli (see Methods; monkey H: median RF eccentricity 6.2 dva, range 5.2-7.1 dva, median RF diameter 0.48 dva, range 0.26-1.88 dva; monkey A: median eccentricity 5.4 dva, range 3.2-8.5 dva, median RF diameter 0.91 dva, range 0.46-2.3 dva). Compared to a surface stimulus of 6 dva diameter, receptive fields had a median proportional diameter of 0.08 (0.04-0.31, monkey H) or 0.15 (0.08-0.38, monkey A). We first studied LFP and MU responses to the presentation of stationary surface stimuli, namely large uniform disks covering the cluster formed by all RFs (6 dva diameter, flashed on and then maintained on screen; Figures 4.2A-B; *Dataset 1*, see Methods). The stimuli did not overlap with the fixation spot. Note that the stimuli were much larger than the RFs of the multi-units, such that they covered a large portion of the multi-units' respective surround regions.

In this chapter, we analyzed differences between chromatic and achromatic surface stimuli (Figures 4.2 and 4.5). In the following Chapter 5, we describe the specific differences among responses to distinct color hues and achromatic stimuli. Note that statistical parameters are largely described in the figure captions.

4.4.1 Characteristics of firing activity and LFP signals in response to uniform surface stimuli

We examined the effect of uniform surface stimuli on LFP power spectra. The presentation of large, chromatic surface stimuli (equiluminant red, green and blue; see Methods) induced prominent, narrow-band gamma oscillations in LFP power spectra (Figure 4.2C-D). These oscillations were clearly visible in the LFP traces (Figure 4.2B). In comparison, gamma-band oscillations were significantly weaker for achromatic surface stimuli (black or white, maximal contrast to background; Figures 4.2C-D).

To test whether V1 spiking activity was synchronized with the induced LFP gamma oscillations, we computed spike-field phase locking spectra (Pairwise Phase Consistency, Vinck et al. 2010b) between MU and LFP activity obtained from nearby but separate sites (Figure 4.2E; see Methods). Spike-field phase-locking spectra for chromatic surface stimuli showed a prominent peak in the gamma-frequency band consistent with the gamma peak in the LFP power spectrum (Figure 4.2D), whereas

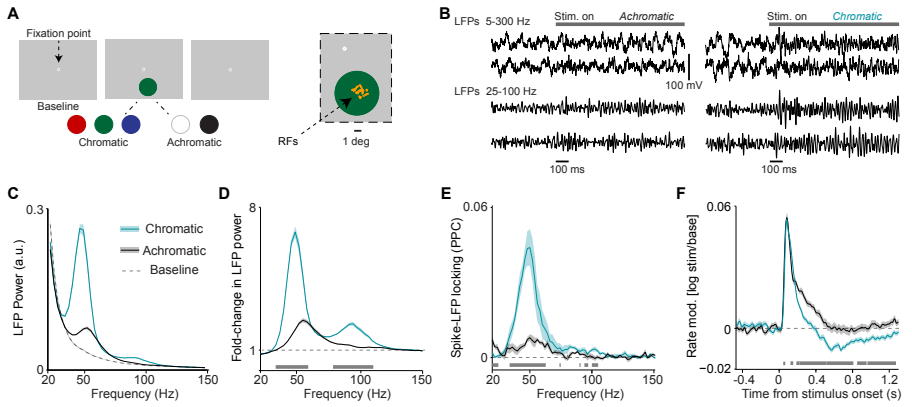

Figure 4.2 — Analysis of LFP and multi-unit activity in response to large, uniform surfaces. (A) Illustration of experimental paradigm with large, 6 dva diameter surfaces (*Dataset 1*, see Methods; n = 5 sessions, 60±0 trials per session for chromatic and 40±0 trials per session for achromatic conditions. Trials numbers were equated by random subselection for statistics.). Trials were self-initiated by fixating on the fixation spot (enlarged for visibility), followed by a baseline period of 0.5-0.8 s with a gray background screen. Surfaces were either chromatic (red, blue, or green) or achromatic (black or white) and presented for 3.3 s, the first 1.3 s of which are analyzed here. Right panel shows the RF locations of analyzed sites in one session. (B) Representative trials of LFP signals for achromatic and chromatic conditions (having gamma power close to the median of the respective condition). (C) Average LFP power spectra for chromatic (turquoise), achromatic (black) and baseline (gray) conditions. LFP power was estimated using Discrete Fourier Transform of non-overlapping epochs of 500 ms, with multi-tapering spectral estimation (± 5 Hz). LFP spectra for all three conditions were normalized to the summed power (> 25 Hz) for the baseline (gray) condition (see Methods). (D) Average change in LFP power, expressed as fold-change, relative to baseline. (E) Average MU-LFP locking, which was estimated using the pairwise phase consistency (PPC, see Methods). (F) Modulation of firing rate relative to baseline, expressed as log(stim/base). (C-F) Shadings indicate standard errors of the means (see Methods). (D-F) Gray bars at bottom of figure indicate significant differences between chromatic and achromatic stimuli, obtained from permutation testing with multiple comparison correction across all frequencies and time points (see Methods).

phase-locking was significantly weaker for achromatic surface stimuli (Figure 4.2E). This demonstrates that gamma-synchronization for chromatic stimuli is not merely observed at the level of synaptic currents within V1, but also at the level of V1 output spikes.

Next, we examined the way in which the presentation of uniform surface stimuli affected MU firing activity. The presentation of chromatic and achromatic surface stimuli induced short-latency onset transients of similar magnitude (Figure 4.2F). However, we observed a stronger decrease in MU firing activity over time during continuous stimulus presentation for chromatic than achromatic surface stimuli, starting around 200 ms after the stimulus onset. Strikingly, for chromatic surface stimuli, MU firing activity fell below baseline levels (Figure 4.2F). Note that in the following chapter, Figure 5.1D, we show that the decrease in MU firing below baseline only occurred for a subset of colors. The reduction in MU firing rates (0.3-1.3 s period) for chromatic as compared to achromatic surface stimuli was observed for 92% of recording sites. Control analyses in which 100 ms data epochs after microsaccades were removed indicate that the late decrease in MU firing was not due to microsaccades. We also found that the LFP results on gamma oscillations were qualitatively unchanged (Figure 4.3).

Our main LFP analyses were performed with smoothing in the frequency domain

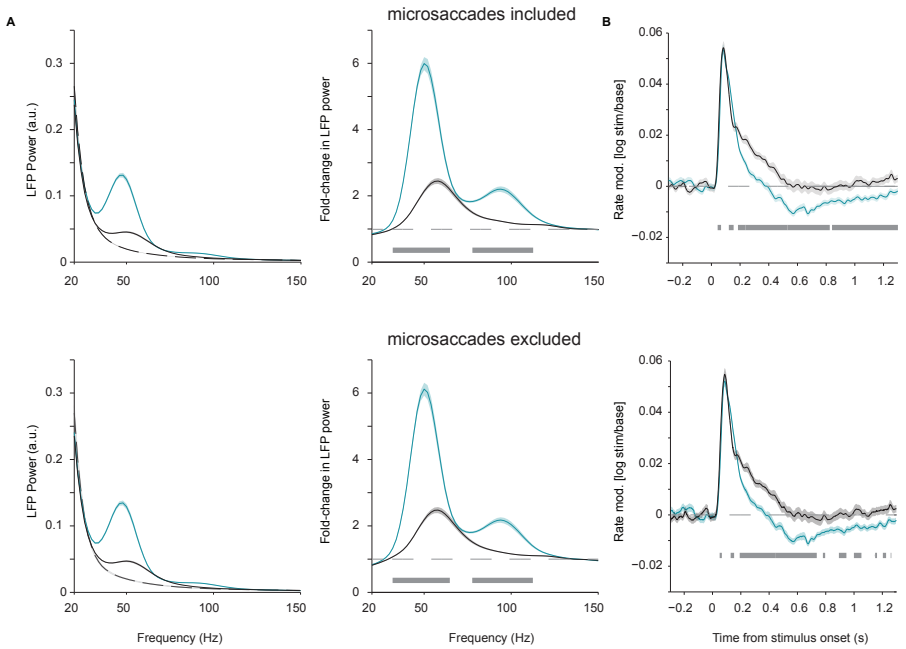

Figure 4.3 — Control analysis for microsaccades. (A) Average LFP power spectra and rate modulation for chromatic (turquoise), achromatic (black) and baseline conditions, separately for data epochs defined irrespective of microsaccades (top) and epochs excluding 100 ms after each microsaccade. Microsaccade detection was based on the algorithm of Engbert and Kliegl (2003) (see Methods). Analysis was based on same dataset as in Figure 4.5. Analyses were performed as for Figure 1, except that the Fourier spectra were computed based on 100 ms epochs that were Hann-tapered and zero-padded to 1 s. (B) Modulation of firing rate relative to baseline, expressed as log(stim/base), separately for data with microsaccade epochs included (top) and excluded (bottom).

and analyzed from 20 Hz onwards (see Methods). In Figure 4.4A, we show log-log LFP spectra down to 2 Hz. Figure 4.4B shows the corresponding relative power spectra with tests for differences between chromatic and achromatic conditions. The effects are specific to the gamma-band. The chromatic-achromatic difference was also highly consistent across sites. For each site and chromatic/achromatic condition, we determined the peak power change in the gamma-frequency range (30-80 Hz) using a polynomial fit (see Methods). Gamma peak power changes were stronger for chromatic than achromatic surface stimuli at 97.8% (45 out of 46) of LFP recording sites (Figure 4.4C).

The LFP power spectra had similar frequency profiles in the two monkeys (i.e. the peaks were well aligned), and the MUs showed similar temporal profiles (Figure 4.4D-I). Therefore, we pooled the data from the two animals in Figure 4.2 and the remaining figures, unless otherwise indicated (the main results are typically also shown per animal).

Note that in the next chapter, section 5.2.2 (Controls for luminance- and cone-contrast), we will describe the results of control experiments in which the chromatic and achromatic stimuli are matched in terms of luminance-contrast and DKL cone-contrast (see Methods) to the full-screen background. We find that the difference in gamma-band oscillations between chromatic and achromatic stimuli is not explained

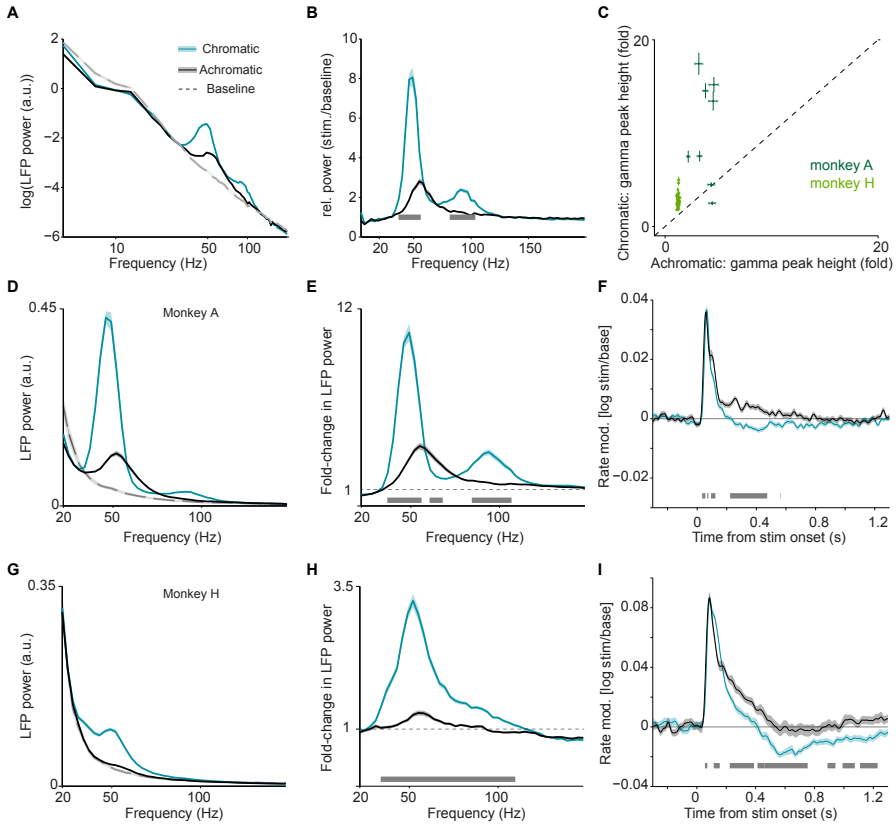

Figure 4.4 — Additional analyses of LFP and multi-unit activity in response to large, uniform surfaces. (A) Average LFP power spectra for chromatic (turquoise), achromatic (black) and baseline (gray) conditions. LFP power was estimated as in Figure 4.2, but with a Hann taper (see Methods). (B) Average change in LFP power, expressed as fold-change, relative to baseline. (C) Scatter-plot for all the LFP recordings sites in two monkeys, showing the gamma-band amplitude (expressed as fold change) in chromatic and achromatic conditions. (D) Average LFP power spectra for chromatic (turquoise), achromatic (black) and baseline (gray) conditions for monkey A, using same estimation settings and normalization for power spectral density as in Figure 4.2. (E) Average change in LFP power, expressed as fold-change, relative to baseline. (F) Modulation of firing rate relative to baseline, expressed as log(stim/base), for monkey A. (G-I) as (D-F), but now for monkey H. (A-I) Shadings and error bars indicate standard errors of the means (see Methods). Gray bars at bottom of figure indicate significance bars, obtained from permutation testing with multiple comparison correction across all frequencies and time points.

by either luminance or cone-contrast to the full-screen background.

These findings demonstrate that large, uniform, chromatic surface stimuli induce low firing activity yet highly gamma-synchronous V1 responses, whereas achromatic surface stimuli induce much weaker gamma-band synchronization but relatively more vigorous firing activity (for further interpretation of this finding, see Discussion).

4.4.2 Dependence of firing activity and LFP signals on stimulus size

The results shown in Figure 4.2 are consistent with the predictability hypothesis (Vinck and Bosman, 2016) outlined in section 4.2.1. Yet, they do not demonstrate directly that

the enhancement in gamma-band synchronization observed for large uniform colored surfaces is due to contextual surround modulation, because we did not manipulate the surround input. Furthermore, it remains unclear whether the observed differences between chromatic and achromatic surfaces can be explained by a difference in contextual surround modulation or other factors like stimulus drive. To address these questions directly, we used a paradigm that varied the stimulus size across trials (Figure 4.5A; see Methods). We selected one site (or a few nearby sites with RF centers within 0.5 dva of the target site) per session and centered the stimulus on the multi-unit's RF, which was previously mapped with moving bars. In each trial a stimulus of a particular size (0.5, 1, 2 or 6 dva diameter) was presented for 600 ms (Figure 4.5A).

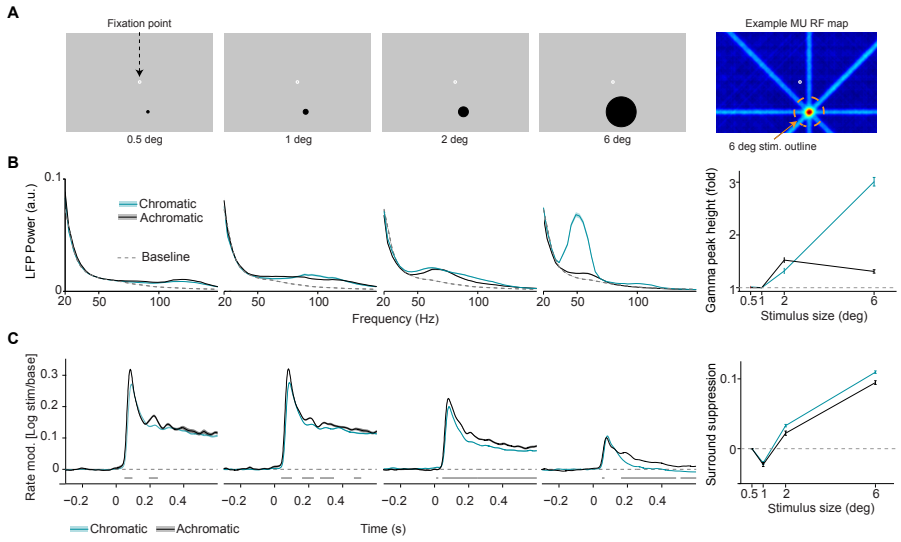

Figure 4.5 — Dependence of LFP power spectra and MU firing activity on surface size. (A) Illustration of experimental paradigm (*Dataset 2*, see Methods; n = 9 sessions, 59.75±0.09/25.64±0.12 trials chromatic vs achromatic trials per condition in each session). Uniform surfaces of four different sizes were presented on a gray background screen. Fixation spot is enlarged for visibility. Right: Receptive field estimated with bar stimuli for a representative target channel, with the outline (orange dashed line) of the largest size stimulus (see Methods). Note that the activation outside the RF is due to the use of large bar stimuli sweeping over the monitor. (B) LFP power spectra for different sizes and chromatic/achromatic conditions. LFP power spectrum estimated and normalized as in Figure 4.2C, but now using 300 ms epochs. Right panel shows the gamma-band amplitude as a function of size, estimated using a polynomial fitting procedure between 30-80 Hz (see Methods). The difference between 6 and 0.5 dva stimuli was significantly larger for chromatic than achromatic condition ($P < 0.05$, bootstrap test, see Methods). (C) Modulation of firing rate relative to baseline, expressed as $\log$(stim/base), for different sizes and chromatic/achromatic conditions. Right panel shows surround suppression, which was defined as the difference in firing rate modulation between the 0.5 degree size and the other sizes.

We first examined how the characteristics of LFP power spectra depended on stimulus size. Analysis of LFP power spectra revealed a strong dependence of gamma power on stimulus size for chromatic stimuli, and by comparison a much weaker dependence for achromatic stimuli (Figure 4.5B). To quantify this size dependence, we determined the gamma peak power between 30-80 Hz (as described for Figure 4.2). For chromatic stimuli, increases in stimulus size resulted in increases in induced gamma peak power as soon as the stimulus also covered the surround (i.e. from 2 dva

onwards, Figure 4.5B). By contrast, for achromatic stimuli, a gamma peak in the 30-80 Hz band emerged from 2 dva stimulus size onwards and showed no further increase with stimulus size.

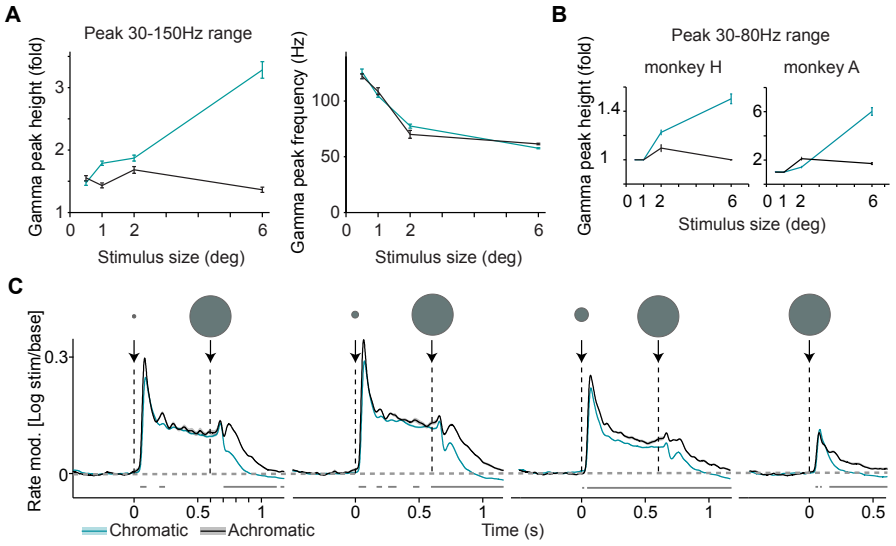

Figure 4.6 — Dependence of LFP power spectra and MU firing activity on surface size: additional analysis. (A) Gamma-band peak amplitude and peak-frequency as a function of size, estimated using a polynomial fitting procedure between 30-150 Hz. A wider range instead of the standard 30-80 Hz range was used here, to also capture peaks >100 Hz, which is far outside the typical range of classical visual gamma. This activity may reflect spike bleed-through, which is beyond the scope of the present study. (B) Average gamma-band peak height in 30-80 Hz range, shown separately for the two monkeys. The difference between 6 and 0.5 dva stimuli was significantly larger for the chromatic than achromatic condition for both animals ($P < 0.05$, bootstrap test). (C) Each trial contained a sequence of two stimuli, either the small stimulus first for 600 ms, or the large stimulus first for 600 ms (see Methods). Here we show the first type of sequence to illustrate the onset of a surround when the stimulus covering the classical RF is already present. Modulation of firing rate relative to baseline, expressed as $\log$(stim/base), for different sizes and chromatic/achromatic conditions. Note rapid firing suppression after onset of the large stimulus following the 0.5 dva stimulus, with a significant difference arising already after ≈100 ms.

We then considered the way in which MU firing was modulated by surround stimulation. We observed that for both achromatic and chromatic stimuli, MU firing activity was highest for 0.5-1 dva stimulus sizes (Figure 4.5C). This was consistent with the estimates obtained from RF mapping and the fact that we centered the presented stimuli on the MUs' estimated RFs. For small stimuli (0.5-1 dva), only the initial transient in MU firing activity showed a difference between chromatic and achromatic conditions, with slightly higher firing activity for achromatic than chromatic stimuli (Figure 4.5C). In contrast, the presentation of a 2 or 6 dva stimulus, increasingly covering the surround, induced strong suppression of MU firing activity as compared to the 0.5 dva stimulus (Figure 4.5C, rightmost panel). This surround suppression was stronger for chromatic than achromatic stimuli (Figure 4.5C).

Furthermore, we analyzed responses during a later period in the trial, when the small stimulus had been presented for 600 ms, and a large (6 dva) surface stimulus of the same color was added for another 600 ms period (4.6C). We found that this addition of

the surround stimulus alone induced a rapid suppression of MU firing activity, which was significantly more pronounced for chromatic than achromatic stimuli (4.6C).

Given the relatively broad increase in >100 Hz LFP power seen in Figure 4.5B, we also determined gamma peak power and peak frequency in a wider range (30-150 Hz). This analysis revealed LFP power peaks >100 Hz for the sizes below 2 dva, and again the strong size dependence for chromatic compared to achromatic stimuli (4.6A; also see 4.6B for an analysis per animal in the original frequency window).

These findings suggest that the relatively strong decrease in firing over time observed for large, chromatic surfaces (Figure 4.2) is at least partially explained by surround suppression. They furthermore indicate that for the small RF stimuli, there are no substantial differences on average between chromatic and achromatic surfaces in terms of MU and LFP responses. Yet, we find a prominent difference in the way chromatic and achromatic stimuli are affected by surround stimulation.

4.4.3 Modulation of firing activity and LFP signals by center-surround predictability

A potential explanation for the results shown in Figure 4.2-4.5 may be the center-surround predictability hypothesis outlined in section 4.2.1 (Vinck and Bosman, 2016). Yet, the employed paradigm used stimuli of different sizes, which may have recruited different neuronal circuits and may also have changed stimulus salience. We therefore used an additional stimulus paradigm in which surround influences were modified, while stimulus size was held constant. Specifically, we created three sets of equally sized stimuli. In one set, the surround was fully predictive of the RF stimulation, because it used a uniform surface (called "uniform" stimulus). In the second set (called "blob mismatch"), the surround was not predictive of the RF stimulation, because the surround stimulus and the 1 dva RF stimulus had different colors (which were physically equiluminant). In the third set (called "annulus mismatch"), the surround had the same color as the RF stimulation, but the two were separated by an annulus ring of a different, physically equiluminant color. This annulus ring had 0.25 dva thickness and an inner diameter of 1 dva).

We found that compared to the uniform surfaces, stimuli with a chromatic (blob or annulus) mismatch had higher MU firing activity (Figure 4.7C). This held true both for the initial transient period and the subsequent sustained response period. At the same time, we observed a marked decrease in the amplitude of LFP gamma oscillations for the chromatic mismatch compared to the uniform surface stimuli (Figures 4.7B-C). This result was consistent across animals (Figure 4.8A).

We further investigated whether this pattern of changes was specific to the sites having RFs near the center of the stimulus. To this end, we examined sites with RFs on the outer uniform regions of the stimulus (with RF centers 1.5-2 dva from the stimulus center; Figure 4.7D). For these sites the MU firing responses did not differ significantly between conditions during the initial transient period (Figure 4.7D). During the later sustained response phase, however, MU firing activity was reduced for the chromatic mismatch stimuli compared to the uniform surface stimulus (Figure 4.7D). Note that whenever the RF center covered a large uniform surface region, either in the uniform stimulus condition or when it covered the surround region of the mismatch stimuli, sustained firing levels were below baseline. This confirmed the respective finding reported in Figure 4.2.

These results suggest that a mismatch between stimuli at the RF center and the surround can dramatically change the surround influence on responses to the center.

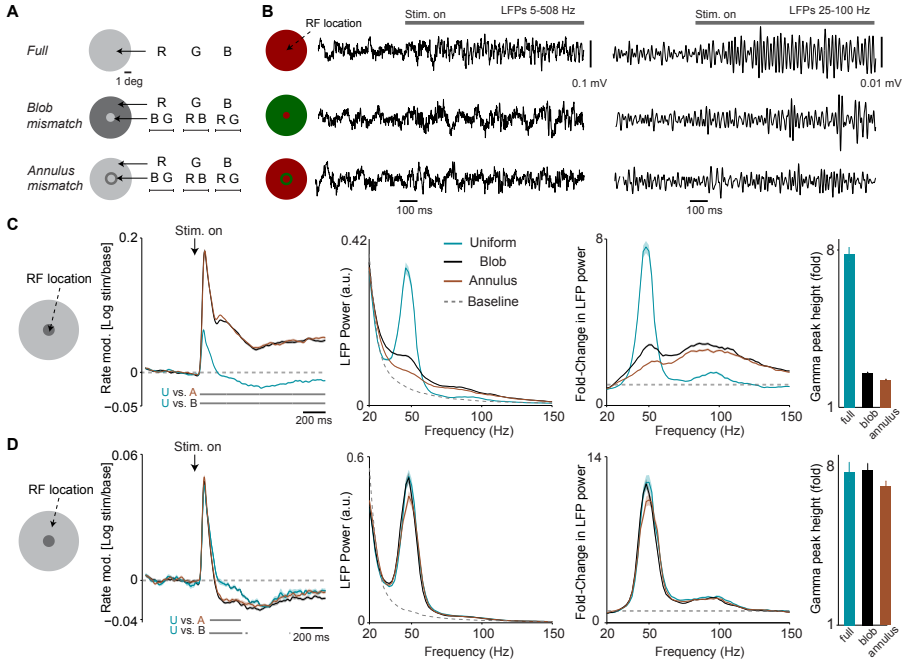

Figure 4.7 — Dependence of LFP power spectra and firing rates on spatial predictability. (A) Illustration of the paradigm (*Dataset 3*, see Methods; n = 8 sessions, 16.26±0.19 trials in each of the 15 conditions per session). We grouped stimuli intro three types. In the "uniform surface" group of conditions, stimuli of 6 dva diameter were presented with either a red, blue or green hue (R B G, equiluminant). In the second "blob mismatch" group, the central 1 dva of the stimulus had a different (equiluminant) color than the rest of the stimulus. In the third "annulus mismatch" group, we presented an annulus ring (of 0.25 dva) of another color on top of the uniform surface (at equiluminant intensity) around the inner 1 dva from the stimulus center. All combinations of hues and stimulus types were presented, yielding a total of 15 individual conditions. (B) Representative LFP traces (having gamma power close to the median of all trials for the respective condition) for the three stimulus types. (C) Analysis for target channels with RFs at the center of the stimulus. Shown from left to right are: (1) The change in MU firing activity relative to baseline expressed as $\log$(stim/base). (2) LFP power spectra for the three stimulus conditions and the baseline. LFP power spectrum estimated and normalized as in Figure 4.2C. (3) The change in LFP power relative to baseline, expressed as a fold-change. (4) The gamma-band amplitude, estimated using a polynomial fitting procedure (see Methods). Gamma-band amplitude was significantly higher for uniform surface than blob and annulus conditions ($P < 0.05$, bootstrap test, see Methods). (D) Same as (C), but now for target channels with RFs between 1.5 and 2 dva from the stimulus center, i.e. close to the central region of the larger, uniform region of the stimulus. Gamma-band amplitude did not significantly differ between conditions (bootstrap test, all $P>0.08$).

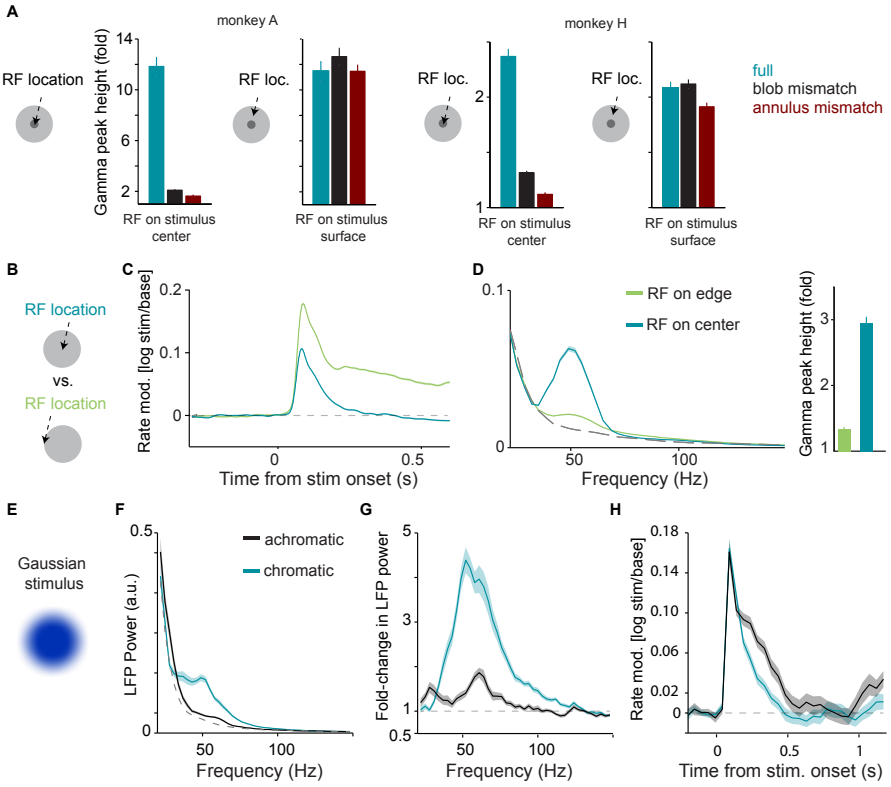

Figure 4.8 — Dependence of LFP power spectra and firing rates on spatial predictability: additional analyses. (A) Comparison of gamma-band power between full surface and mismatch conditions, separately for the two monkeys. For both monkeys, gamma-band amplitude was significantly higher for uniform surface than blob and annulus conditions when the RF was at the center of the surface stimulus ($P < 0.05$, bootstrap test, see Methods). (B) Illustration of paradigm (*Dataset 2*). Uniform surfaces were either centered on the unit's RF, or the edge of the surface was centered on the unit's RF. (C) Modulation of firing rate relative to baseline, expressed as log(stim/base). (D) Average LFP power spectra, using the same analysis time window and spectral estimation parameters as in Figure 2 of main text, comparing "RF-on-center" and "RF-on-edge" conditions. Dashed gray line corresponds to baseline (gray background screen). Right: Gamma-band amplitude (expressed as fold-change) for the two conditions. Gamma-band amplitude was significantly higher for "RF-on-center" condition ($P < 0.05$, bootstrap test). (E)-(H): Single session (from monkey H) illustrating responses to Gaussian surface stimuli. Stimuli were otherwise the same as *Dataset 1*. Gamma oscillations were not abolished by removal of the sharp stimulus edge. (E) Example Gaussian stimulus that had a blurred edge. (F) Average LFP power spectra for chromatic (turquoise), achromatic (black) and baseline (gray) conditions, computed as in Figure 4.2. (G) Average change in LFP power, expressed as fold-change, relative to baseline. (H) Modulation of firing rate relative to baseline, expressed as log(stim/base), for monkey A.

We wondered whether the surround influence on gamma oscillations originates from the uniform surface region or rather from the edge of the surface. To this end, we analyzed sessions in which we compared two sets of trials: First, trials with a full surface stimulus centered on a site's RF ("RF-on-center" condition; Figure 4.8B). Second, trials with a full surface stimulus positioned such that its edge fell into the RF center, i.e. with the surface shifted by 3 dva horizontally ("RF-on-edge" condition; Figure 4.8B. We found that the amplitude of gamma oscillations was significantly higher at the center ("RF-on-center") than at the edge of the surface stimulus ("RF-on-edge"), whereas the opposite was observed for MU firing activity (Figure 4.8B).

In one session (monkey H), we also showed disk stimuli that had their edge blurred with a Gaussian (2.5 dva size, 1 dva standard deviation). There were clear gamma responses also in this case (Figure 4.8C).

Together, these results indicate that for colored surfaces the amplitude of gamma-band oscillations is commensurate with the "chromatic" predictability among visual inputs in space, and that gamma-band oscillations are not a mere consequence of input drive to a larger cortical region. Furthermore, these results suggest that gamma strength can be dissociated from stimulus salience, because the chromatic mismatch condition provided a highly salient stimulus in the RF, but resulted in weaker gamma.

4.5 Discussion

4.5.1 Differences between chromatic and achromatic surfaces

We found large differences in V1 gamma-band activity and firing responses between chromatic and achromatic stimuli. A previous voltage-sensitive dye imaging study has shown differences in V1 responses to chromatic compared to achromatic surfaces (Zweig et al., 2015), but did not (and could not) consider gamma or rate responses due to the imaging method employed. This study found that, specifically for achromatic but not chromatic surfaces, the dynamics of V1 responses was in line with a fill-in process (from the edge of the surface toward the center). This is in line with the notion that chromatic surface responses are dominated by single-opponent LGN inputs to V1 and responses of V1 neurons with RFs within the uniform surface (Zweig et al., 2015; Livingstone and Hubel, 1984; Shapley and Hawken, 2011). Indeed, some neurons with chromatic opponencies in LGN and V1 show rate increases for large chromatic surfaces, but do not respond to large achromatic surfaces (Ts'o and Gilbert, 1988). Therefore, it is possible that LGN and/or layer 4 drive within V1 is stronger for chromatic than achromatic stimuli. This view is also supported by the strong source of drive for chromatic compared to achromatic stimuli in LFP signals reported by Chen et al. (2007). On the other hand, some neurons will be driven for achromatic stimuli when there is a temporal luminance change (Xing et al., 2010). The fill-in process from the edge described above is also a form of (indirect) drive (Zweig et al., 2015). Therefore, chromatic and achromatic stimuli likely differ in input drive, but could also be processed in fundamentally different ways.

fMRI signals tend to be stronger in response to chromatic compared to achromatic stimuli (for review, see Schluppeck and Engel 2002, Shapley and Hawken 2011). Given the complexity of the biophysics of the fMRI signal, this could reflect stronger rate responses in the input layers, but also the strong gamma responses reported here. Several studies have reported that the fMRI signal can correlate not only with spiking activity, but also (and sometimes more strongly), with gamma-band activity (Logothetis and Wandell 2004; Ekstrom 2010; Maier et al. 2008; Thomsen et al. 2004; Viswanathan and Freeman 2007; Nir et al. 2007; Scheeringa et al. 2016, for visual gamma in particular Bartolo et al. 2011; Niessing et al. 2005).

Here, we observed stronger gamma-band oscillations yet more suppressed firing responses for chromatic surfaces. A possible interpretation of this is the direct activation of neurons with RFs in the "predictable" part of the surface for chromatic surfaces, which is subsequently "explained away" and converted into gamma-band activity through network interactions. This suggests that one could think of color as a "feature" (like orientation or motion), and that this feature space is used for prediction in space. In cases where the surrounding spatial contexts correctly predicts

the central stimulus, there will be strong gamma. In contrast, predictability based on luminance, as given for achromatic (and chromatic) stimuli, may not be treated as a "feature" by cortex. Luminance by itself, in the absence of color information, may not be sufficient for inducing strong gamma-band synchronization. Speculatively, this may be related to the fact that illumination can generate large luminance "surfaces" that are object-unrelated, whereas hue continuity may be informative about visual objects. One should note that, prominent gamma-band synchronization *can* be generated in response to achromatic stimuli. This tends to occur when they have spatial structure (orientation, frequency, phase) that is highly predictable, e.g. bars and gratings (Gieselmann and Thiele, 2008; Gray et al., 1989; Chalk et al., 2010; Gail, 2000; Singer, 2018).

4.5.2 Drive and predictability as ingredients for gamma-band synchrony

The present data suggest that V1 gamma-band activity requires both sufficient drive and spatial predictability. The necessity of drive is indicated by the weak gamma responses to achromatic stimuli, and is supported by previous findings that low-contrast stimuli generate weak gamma-band activity (e.g. Ray and Maunsell 2010; Roberts et al. 2013, Hadjipapas, Lowet, Roberts, *Peter* and De Weerd, (2015)). In the following chapter, we will see that prolonged adaptation also results in weak gamma responses (Figure 5.7).

Note that in several experiments, V1 gamma-band amplitudes and firing rate (our best available proxy for drive) show opposite effects. In particular, Figure 4.7 shows strong gamma oscillation decreases along with large rate increases for mismatch compared to uniform conditions. Similarly, gamma-band activity was stronger at the stimulus center, but rate responses were stronger at the edge of chromatic surfaces (Figure 4.8B). These strong responses to edges may derive from double-opponent cells that are driven by chromatic edges and are found in superficial layers (Shapley and Hawken, 2011; Johnson et al., 2008; Friedman et al., 2003; Shapley and Hawken, 2002). Our recordings were biased toward supercifial layers (see Methods). Notably, strong gamma generators appear to exist in these layers (see General Introduction, Figure 1.4). Taken together, the current experiments show strong gamma and weak firing at predictable locations in visual space, and the opposite at unpredictable locations. This is in line with the "predictability hypothesis" outlined in Relationship between gamma, surround modulation and sparse coding. The discussion of both functional and mechanistic interpretations will be continued in the General Discussion (Chapter 6).

4.6 Chapter summary

This chapter demonstrated that V1 responses to chromatic and achromatic surfaces are differentially modulated by spatial context.

1. Compared to achromatic surfaces, chromatic surfaces induced strong gamma-band synchronization of both LFP and multi-unit responses (Figures 4.2 and 4.5).
2. Initial firing transients between chromatic and achromatic surfaces were similar (Figure 4.2 and 4.5), yet in the sustained stimulation period, firing responses to chromatic surfaces were lower. This was a result of a stronger decrease in firing over time, due to an increase in surround suppression (Figures 4.2 and 4.5).
3. For colored stimuli, those with a "chromatic mismatch" between the stimulus covering the RF center and the surrounding surface evoked high firing activity, yet a very prominent reduction of gamma-band activity compared to uniform, "predictable" surfaces (Figure 4.7). This effect was localized only to the multiunit sites whose RFs covered the mismatch.

Chapter 5

Stimulus-hue and background-hue dependence of gamma

5.1 Introduction

5.1.1 Hue

In the previous chapter, we found that chromatic, but not achromatic surfaces induce strong gamma-band synchronization yet low firing rates when they 1) cover both the RF and its surround and 2) have a predictable center-surround relationship. In these analyses, all chromatic conditions were pooled irrespective of the hue, because there appeared to be a qualitative difference between responses to black or white compared to colored stimuli. However, it is reasonable to assume that different color hues may elicit different responses. Color hues are a central feature for object recognition and visual search, and have high ethological relevance in social interactions and foraging (Santos et al., 2001; Waitt et al., 2006; Gerald et al., 2007; Bichot et al., 2005; D'Zmura, 1991; Corso et al., 2016). In particular, many primate species, including macaques and irrespective of their own sex, have a viewing preference for reddish/pink facial and/or genital regions (Waitt et al., 2006; Corso et al., 2016; Dubuc et al., 2014). Red-green differences as established by trichromacy are also relevant for food retrieval (Melin et al., 2017).

Several studies, using fMRI or extracellular recordings, have studied responses to different hues in V1. (For an introduction to color processing in V1, see Chapter 1 section Color in V1). Different V1 cells exhibit a large variety of tuning to different hues, with response peaks at different hues, different degrees of selectivity for a particular hue, etc. (Shapley and Hawken, 2011; Solomon and Lennie, 2007; Livingstone and Hubel, 1984; Wachtler et al., 2003; Johnson et al., 2001). Responses are more complex than in the input structure, the LGN (Shapley and Hawken, 2011; Livingstone and Hubel, 1984; Wachtler et al., 2003; De Valois et al., 2000), and seem to undo some of the strong bias toward more frequent red-green than blue-yellow responses observed in its input (De Valois et al., 2000; Wachtler et al., 2003). There can be nonlinear interactions between hue and luminance (Horwitz et al., 2004). V1 single cells cover all possible hues with their response peaks, but they do not do so uniformly. Wachtler et al. (2003), using surface stimuli, report several over-represented hues, among them yellow-orange and a greenish hue that is maximally off the L-M axis. De Valois et al. (2000) find other hues to be over-represented using other stimuli. Overall, the literature agrees that hues representation is non-uniform, also assessed with fMRI (Lafer-Sousa et al., 2012), but substantially more continuous than in LGN, where responses are strongly clustered along the red-green and blue-yellow axes (Wachtler

et al., 2003; De Valois et al., 2000). Despite the importance of the color domain in vision, hue differences have rarely been studied with respect to gamma synchronization (Rols et al., 2001; Shirhatti and Ray, 2018). Rols and colleagues (2001), using electrocorticography, reported stronger responses to red than green stimuli across a range of luminances, with intermediate values for yellow stimuli (their Figure 6c). Similarly, Shirhatti and Ray (2018), working independently from the results presented here, found strong gamma LFP power for full-screen uniform chromatic stimulation, particularly for red, but also for blue hues.

5.1.2 Contextual dependence of chromatic responses

Local spatial context plays a strong role in vision (see previous chapter and Chapter 1 section 1.4.2). So does temporal context on various timescales (see Chapter 2). More broadly, the overall context matters. For color, this is particularly true in terms of background illumination. For example, we are able to discriminate a red from a green apple[1] under various lighting conditions, to the degree that the red apple in one lighting condition may have the same objective spectral properties as the green apple under different illumination. How is this possible? The answer is that the color percept is based on the current context. V1 seems to be involved in this computation (Wachtler et al., 2003) and related perceptual consequences. A more extreme example of contextual effects on perception are the visual aftereffects after adaptation through prolonged viewing, as decribed in the Introduction, section 1.4.1. The visual system will therefore show adaptation specific to the background. In the current study, this effect can be used as a tool to explore the robustness of differences in responses to hues across contexts, and illuminate its underlying mechanisms. After considering differences between colors under standard viewing conditions, we will therefore explore the effects of different contexts.

5.2 Results

5.2.1 Differences in gamma-band activity between color hues

Preliminary remarks: The methods for this chapter are identical to the ones in the previous chapter, see Chapter 4 Methods. Similarly to the previous chapter, results are largely taken over ***verbatim*** *from the respective publication, Peter et al. (2019), see also Acknowlegdements for Chapters 4 and 5.*

The results in the previous chapter showed prominent differences between chromatic and achromatic surfaces in terms of gamma-band synchronization. The respective analyses pooled different chromatic conditions (equiluminant red, green and blue) together. However, there may exist further differences in gamma-band synchronization within the chromatic conditions, i.e. between different hues. To investigate this we used two types of stimulus sets, which were presented in separate sessions. In the first stimulus set (Figure 5.1A), we presented each surface color at its maximum possible luminance level (given the limits of the employed monitor), and sampled from the entire spectrum of hues available with the monitor (see 4.3 Methods). In the second stimulus set, we presented surface stimuli with different color hues at equated luminance levels (Figure 5.1B-C).

[1] unless we happen to be colorblind

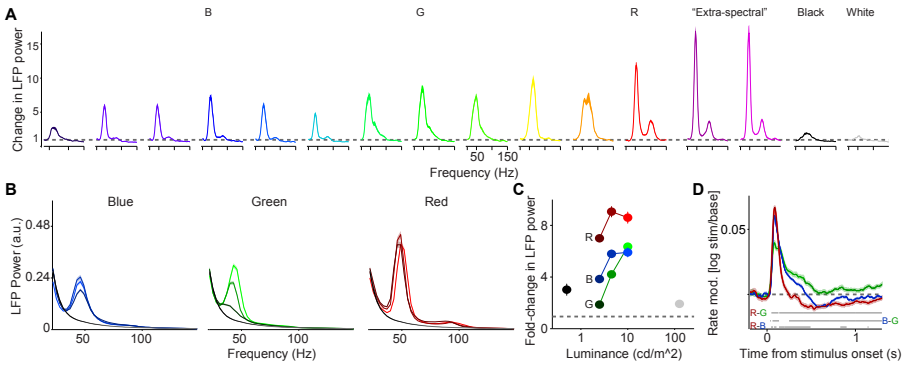

Figure 5.1 — Dependence of LFP power spectra and MU firing activity on surface hue and luminance. (A) We presented uniform surfaces of 6 dva diameter at maximum possible luminance levels, sampling from the available spectrum of wavelengths (*Dataset 4*, see Chapter 4 Methods; n = 5 sessions, 15.24±0.1 trials per condition in each session). In addition, we presented black and white surfaces. Shown is the change in LFP power relative to the baseline (gray screen), expressed as a fold-change. (B) Three hues (red, green and blue) were presented at three different luminance levels (approximately 2.5, 5 and 10 cd/m^2, *Dataset 1*, see Chapter 4 Methods; n = 5 sessions, 19.8±0.45 trials per condition in each session). Shown are LFP power spectra, estimated and normalized as in Figure 4.2C. (C) Average change in LFP power, expressed as a fold-change, relative to the baseline (gray screen). The dependence of gamma amplitude on stimulus luminance was greater for G than for R or B (difference between high versus low: $P < 0.05$, bootstrap test). Gamma oscillations amplitude R>B or G across all three luminance conditions, and B>G for low and intermediate luminance conditions ($P < 0.05$, bootstrap test). (D) Modulation of firing rate relative to baseline, expressed as $\log$(stim/base). Horizontal bars at bottom of panel represent significant differences between stimuli at $P < 0.05$ (permutation test, multiple comparison corrected for time bins). (B-D) Color hues were adjusted for better discriminability, panel A of Figure 5.3 shows actual hues.

Using the first stimulus set, we found that gamma-band LFP oscillations were reliably induced across the entire spectrum of hues (Figure 5.1A, Figure 5.2AB, see also supplementary Table 5.5 for all luminance and CIE values). In addition, we found that gamma-band responses were reliably induced by surfaces with "extra-spectral" colors, i.e. colors resulting from a mixture of blue and red primaries (Figure 5.1A), as well as brownish hues (Figure 5.2B).

Furthermore, we replicated our finding that gamma oscillations were relatively weak for both black and white surface stimuli as compared to all colored surfaces (Figure 5.1A). In one monkey (A), we found that gamma-band activity was stronger for black than for white stimuli, consistent with previous results showing stronger firing rate responses to black than white stimuli (Xing et al., 2010; Yeh et al., 2009). However a trend in the opposite direction was observed for monkey H (Figure 5.3D).

For the first stimulus set (Figure 5.1A), the different colors were presented at their maximum possible luminance levels, which might confound the effects of hue and luminance. We therefore used a second stimulus set in which we presented surface stimuli with different color hues at three levels of equal physical luminance, i.e. different color values (Figures 5.1B-C). For all three hues, gamma amplitudes were greater for the highest compared to the lowest luminance condition ($P < 0.05$, bootstrap test, see Chapter 4 Methods; Figure 5.1B-C). The dependence of gamma amplitude on stimulus luminance was greater for green than for red or blue surface stimuli (difference between high versus low, Figure 5.1B-C). Gamma oscillations had a higher amplitude for red than for blue or green surface stimuli across all three luminance conditions

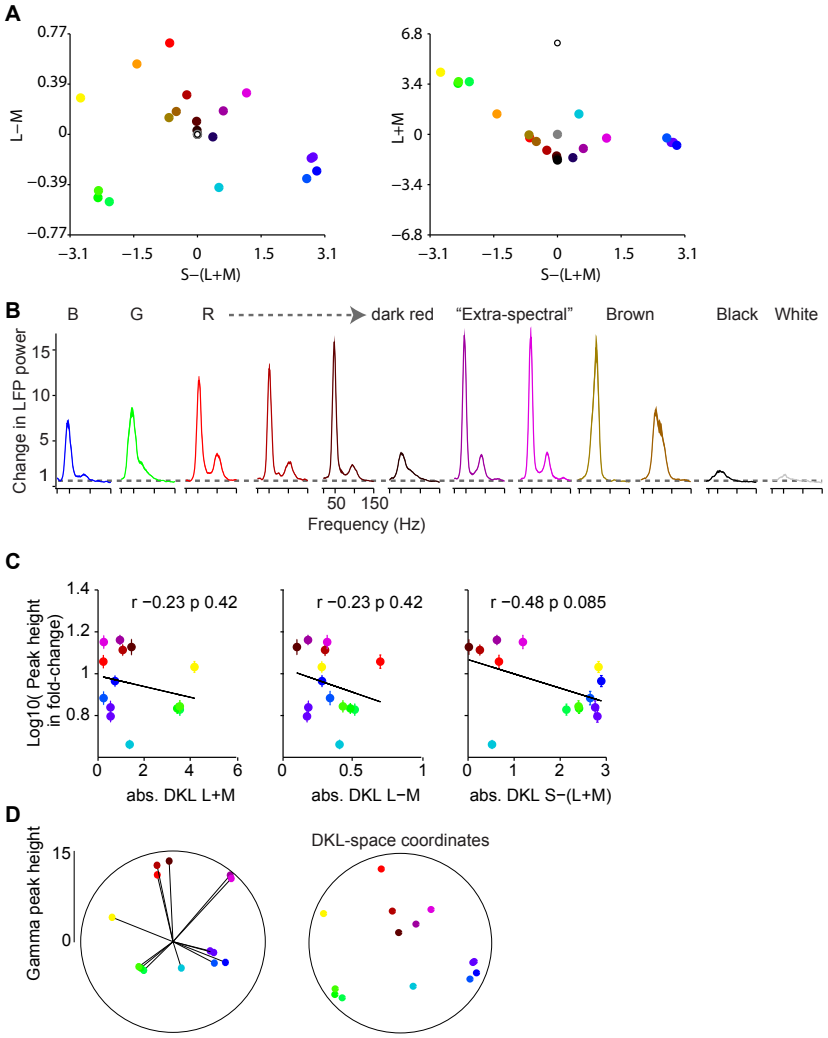

Figure 5.2 — Dependence of gamma LFP power on the surface hue: additional analyses in relation to Figure 5.1. (A) Representation of hues used for Figure 5.1A in DKL space. (B) LFP power spectra as in Figure 5.1A, for additional hues. The pure blue, green and red stimuli are shown for reference, followed by red surfaces of decreasing luminance and the brown surface stimuli. See Table 5.5 for CIE and luminance values. (C) Correlations of gamma-band amplitude ($\log$ of peak height estimate) with absolute cone contrasts in DKL space were not significant, with trends for negative correlation values. Log-transformed values were used to reduce the effects of outliers on correlations. The pattern was similar across monkeys, with monkey A showing a significant, negative correlation of absolute S-(L+M) value to peak height ($P = 0.039$ corrected for multiple comparisons within but not across animals). (D) Gamma peak height as a function of angle in the DKL plane. Angles were obtained by normalizing the L-M and S-(L+M) axes to the maximum absolute value for each respective axis. The resultant vector length was at an angle of 73.95 degrees (where 90 degrees is defined as maximum L-M contrast and no S-(L+M) contrast) and significantly clustered (resultant length = 0.33, $P < 0.001$, permutation test). The panel on the right shows the DKL space coordinates of the colors used in the left panel.

(Figure 5.1B-C), whereas gamma amplitude was higher for blue than green surface stimuli for low and intermediate luminance conditions (Figure 5.1B-C). Another difference between the hues was that the gamma peak had a significantly lower frequency for green compared to red or blue surface stimuli ($P < 0.05$, bootstrap test; Figures 5.1B-C and 5.3). The results of these analyses were consistent across both monkeys (Figure 5.3).

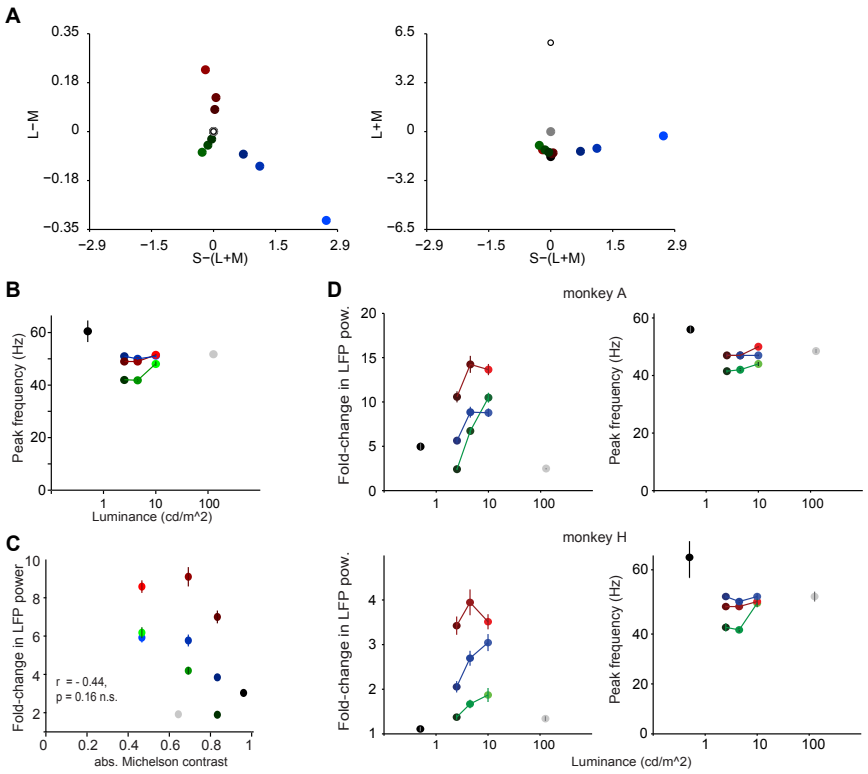

Figure 5.3 — Dependence of gamma LFP power on the surface hue: further analyses in relation to Figure 5.1B. (A) Representation of hues used for Figure 5.1B in DKL space (see Chapter 4 Methods). (B) Peak frequency estimates (Hz) based on cross-validated fitting procedure (see Methods) for the surface stimuli of Figure 5.1B. (C) A regression of absolute luminance contrast (Michelson contrast) against gamma peak height (fold-change) showed no significant relationship (p=0.23). Note that since there was relatively little gamma power for achromatic, high-contrast stimuli, if anything there would be a negative relationship between luminance contrast and gamma power, the very opposite of findings about gamma power for achromatic gratings. Also note that for red stimuli, gamma power appears to follow a U-shape with decreasing luminance and increasing contrast (see also Figure 5.1A, 5.2). (D) Gamma power (fold change) and peak frequency estimates (Hz) for the data of Figure 5.1B, per animal. For both animals, the brightest stimuli showed stronger gamma-responses than the darkest stimuli. The difference in fold-change gamma amplitude as a function of luminance was greater for green than for red for both animals and significantly greater than blue in monkey A. Moreover, red responses across luminances were stronger than green or blue responses, with the exception of the brightest blue responses in monkey H. (B-D) Color hues were adjusted for better discriminability given the small dot sizes, (A) shows actual hues.

Given the relationships between MU firing activity and LFP gamma-band oscilla-

tions shown in the previous chapter in Figures 4.2-4.7, we asked how these differences in LFP gamma oscillations were related to changes in firing activity. During the initial transient, MU firing activity was higher for red and blue rather than for green surface stimuli, with slightly stronger responses for red than blue surface stimuli (Figure 5.1D). Yet, we found that the post-transient decrease in MU firing activity over time was particularly pronounced for red and particularly weak for green stimuli (Figure 5.1D). In agreement with the data shown in Figure 4.2, we observed that MU firing activity fell below baseline levels for red and blue surface stimuli (Figure 5.1D).

Together, these results indicate that surfaces of all color hues tend to induce gamma-band oscillations with a higher amplitude than achromatic surfaces, and that the amplitude of gamma oscillations is relatively high for red surfaces.

5.2.2 Controls for luminance- and cone-contrast

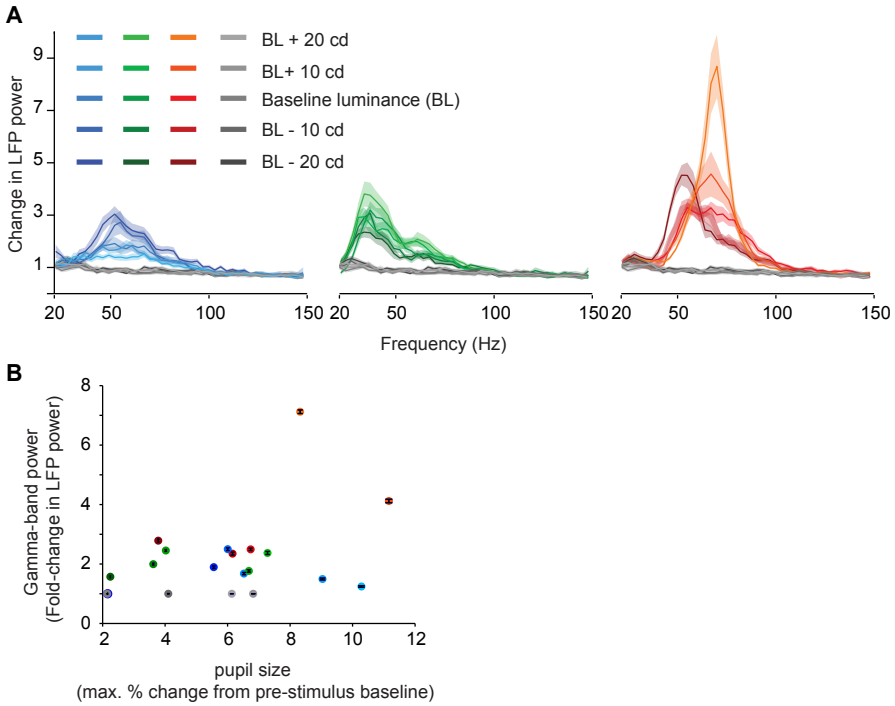

Figure 5.4 — Dependence of gamma LFP power on the surface hue: control experiment for luminance-contrast and pupil responses. (A) Single session from monkey H with luminance match between achromatic and chromatic surface stimuli. Each color, including gray, was shown at 5 luminance levels, namely equiluminant to the background gray, or in two steps of 10cd/m^2 brighter or darker than the background. Blue and red hues at higher luminances were obtained by adding luminance from the green channel to maximal blue or red output, respectively. Note that achromatic responses are weaker than chromatic responses also in this luminance-matched case. (B) Dependence of gamma-band power on pupil responses. Pupil responses represent the maximal percent change from baseline during the stimulus period. The correlation over all hues was not significant ($r = 0.38$, $p = 0.1$). Note that gamma responses for achromatic stimuli were poor regardless of the degree of pupil change, which spanned a large part of the range occupied by the chromatic stimuli. Error bars are included but typically smaller than the dot size.

In the analyses above and the previous chapter, we observed a strong difference

in gamma-band power between chromatic and achromatic surfaces. We performed several control analyses and experiments to investigate whether this observed difference was explained by differences in DKL cone contrast or luminance contrast between chromatic and achromatic surfaces. Firstly, a linear regression of gamma peak height against absolute Michelson luminance contrast ((luminance stimulus - luminance baseline) / (luminance stimulus + luminance baseline)) across the surface stimuli shown in Figure 5.1B showed no significant relationship (r = -0.44, P = 0.16, F-test, Figure 5.3; note that the relationship, if any, was negative). Secondly, in an additional control experiment, we directly matched the luminance (and thereby luminance-contrast) of the achromatic and chromatic stimuli across 5 brightness values, including the full-screen background brightness and two steps of positive and negative contrast. We found that achromatic gamma-responses were much weaker than chromatic gamma-responses regardless of overall luminance level, also under these matched conditions (Figure 5.4A). We additionally used this experiment as a control for the effect of pupil size (see Chapter 4 Methods) on gamma-band amplitudes (Figure 5.4B). Note that gamma responses for achromatic stimuli were weak regardless of the degree of pupil change.

Thirdly, in another experiment, which is part of the data shown in Figure 5.7, we matched cone-contrasts between chromatic and achromatic stimuli. Specifically, we compared gamma-responses to a colored surface on an achromatic full-screen background with its "inverse", gamma responses to a corresponding achromatic surface on a chromatic full-screen background (e.g. red on a gray background versus gray on a red background). These comparisons keep the changes in cone-activation relative to the background the same. Note that this does not mean that the cone-contrasts are matched in the DKL space, because this space contains an additional normalization step, which incorporates the extent to which the full-screen background itself activates the different cones. Nevertheless, although only the non-normalized changes in cone-contrasts are matched, it can be seen that e.g. the white stimuli have very strong DKL cone-contrast to the chromatic full-screen backgrounds along the L-M and S-(L+M) axes. This cone-contrast for white surfaces on chromatic backgrounds exceeds that of chromatic stimuli on the white background (see supplemental Figure 5.11). Our analyses reveal that for each tested color (except for red on a black surface), gamma was much stronger for chromatic than achromatic surfaces of matched cone contrast (Figure 5.7F). Together, these data indicate that the difference in gamma-band response between chromatic and achromatic surfaces was not due to luminance- or DKL cone-contrast relative to the full-screen background.

In the previous section, colored stimuli were either presented at maximum brightness or presented at the same physical luminance. This leaves open the possibility that differences observed between hues are due to differences in cone contrast (see Chapter 4 Methods). We performed additional experiments in which colored surfaces were matched in terms of DKL space coordinates in units of Weber cone contrast (see Chapter 4 Methods, Figure 5.5). These coordinates were the L-M (red-green opponency), S-(L+M) (blue-yellow opponency) and L+M (luminance) cone-contrasts relative to the gray full-screen background. In the first experiment (Figure 5.5A), we selected three luminance steps (L+M cone-contrast was -0.25, 0, or +0.25). For each luminance step, we then took an equal step in the positive and negative L-M direction. This step was taken as the maximum possible step for which the magnitudes were equal in both directions. Similarly, we took a step of equal magnitude in the positive or negative S-(L+M) direction. In the second experiment (Figure 5.5B), we sampled from 8 different angles in the DKL plane at an equiluminant level to the gray back-

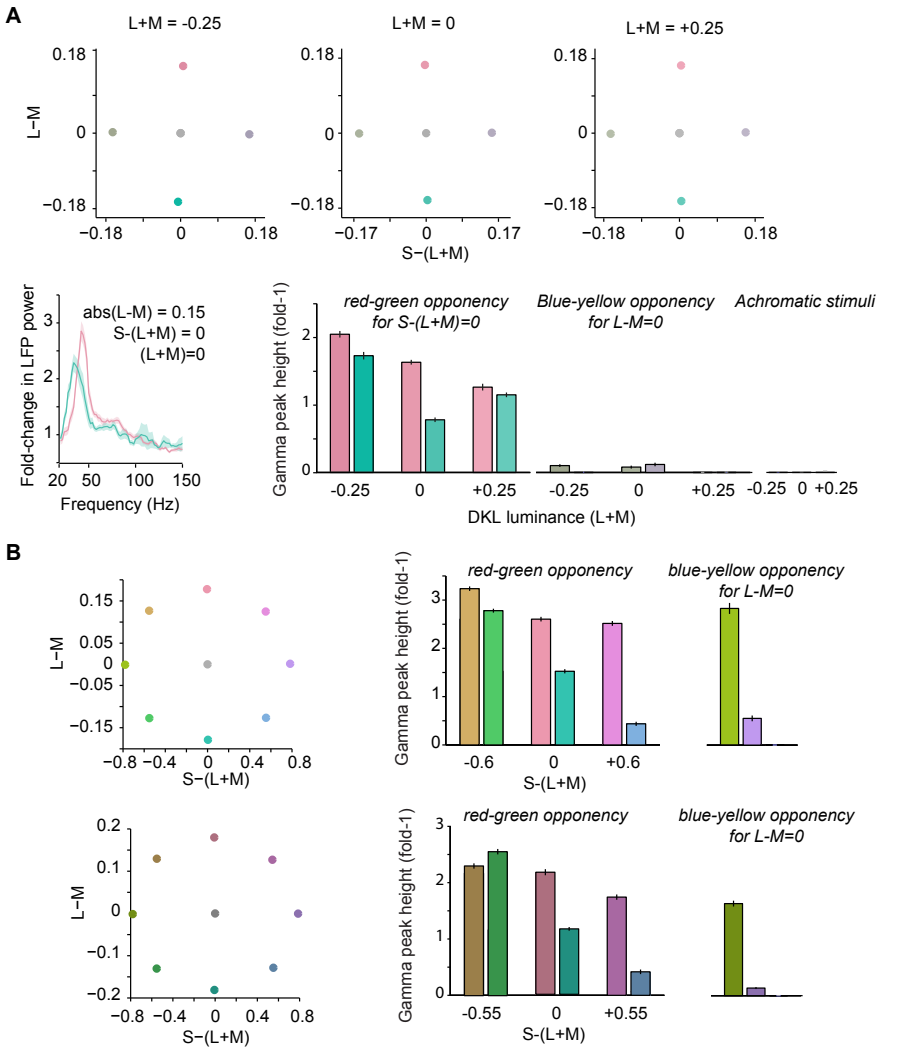

Figure 5.5 — Gamma-band power for stimuli defined on equiluminant DKL planes, control experiment related to Figure 5.1. (A) The DKL-space plots at the top illustrate the employed colors. They were designed to present stimuli with equal cone-contrast in L-M (red-green opponency) and S-(L+M) (blue-yellow opponency) directions, separately for three luminance (L+M) steps of -0.25, 0 and +0.25 relative to the gray background (having RGB values of 175, 175, 175) (N=1 session in Monkey H). These luminance intensities corresponded to 54.7± 7.9 cd/m^2. The spectra and bar plots below show the resulting gamma-band responses. Gamma-band oscillations were stronger for reddish than greenish stimuli and very weak or non-detectable for the blue/yellow and achromatic stimuli. Note however that the blue/yellow components were perceptually very faint at these S-(L+M) values that were matched to the L-M cone-contrast values. In (B) we therefore used a normalization of the S-(L+M) axis to the most negative S-(L+M) value obtainable (along the L-M=0 axis). (B) Top: Gamma-band activity for stimuli with equal cone-contrast in positive and negative L-M directions, and equal cone-contrast in positive and negative S-(L+M) directions (N=1 session in monkey H). Stimuli were equiluminant relative to the gray background, which was the same as in (A). Gamma was stronger for reddish than greenish hues. In addition, with a stronger S-(L+M) cone-contrast as compared to panel (A), prominent gamma-band activity was now generated in the absence of a cone-contrast in the L-M direction. This gamma-band activity was stronger in the yellow (positive S-(L+M) values) than the blue direction. Bottom: Same as top, but now for a darker gray background (the same as used for Figure 4.2 in the previous chapter) and darker surface stimuli (of matched luminance to background, N = 1 session in Monkey H).

ground. Note that this yields stimuli that are highly desaturated as compared to the stimuli shown in Figure 5.1.

The data from these control experiments show that gamma-band responses were stronger for reddish than greenish hues. This is consistent with the results shown in Figure 5.1A (see Figure 5.2D). Achromatic stimuli did not induce detectable gamma-band response peaks, neither in the positive nor negative luminance-contrast (L+M) step, consistent with the findings of the previous chapter, Figures 4.2-4.5, and also Figure 5.1 and 5.4A. Furthermore, these data suggest that increasing S-(L+M) cone-contrast (blue-yellow opponency) independently (or in absence) of L-M cone-contrast also boosts gamma oscillations (Figure 5.5B). This contrasts with the study of Shirhatti and Ray (2018), who reported no correlation of strength of gamma-band power with blue-yellow contrast in their stimulus set (see Discussion). These results provide further support for the notion that gamma-band oscillations for uniform surfaces are mediated by color-opponency signals. They indicate that hue (i.e. the angle in the DKL-plane) itself is a determinant of gamma-band oscillation strength, and that the dependence on hue is not explained by the magnitude of cone-contrasts as a confounding influence. This is also consistent with the finding that the magnitudes of the DKL cone-contrasts for the chromatic surface stimuli shown in Figure 5.1A are not significantly correlated with gamma-band power (Figure 5.2C). Note that this lack of a correlation might be due to the use of stimuli with largely very high brightness and therefore cone contrasts, such that the effects of hue differences dominate. These experiments were not designed to detect the effect of variations of cone-contrasts within a particular hue (see Discussion).

5.2.3 Temporal evolution of gamma-band responses

The observed differences in gamma oscillations between the color hues (Figure 5.1) might reflect a static and context-independent property of visual cortex to respond differently to distinct hues. Yet, the continuous presentation of a uniform surface stimulus for the duration of an entire trial (several seconds) likely induces substantial adaptation at many levels of the nervous system. We thus wondered whether different hues might adapt at different rates. To address this, we examined the temporal evolution of LFP power spectra over a longer time period, i.e. up to 3 s after stimulus onset.

Time-frequency representations showed that qualitative differences between hues and between luminance levels tended to be relatively stable over time (Figure 5.6A). However, we found that the amplitude of gamma-oscillations decreased more rapidly over time for green than for blue or red surface stimuli (Figure 5.6A-B). This result held true also for both animals individually (monkey H/A red change over time -0.07±0.05/-0.12±0.04, green -0.51±0.06/-0.36±0.04, blue -0.23±0.05/-0.10±0.04). The main effect of decrease with time, as well as stronger decreases for green compared to both red and blue, were significant in both animals individually (all $P < 0.05$ corrected for multiple comparisons). This suggests that there may be differences in the time course and strength of adaptation between color hues, specifically stronger adaptation for green surface stimuli.

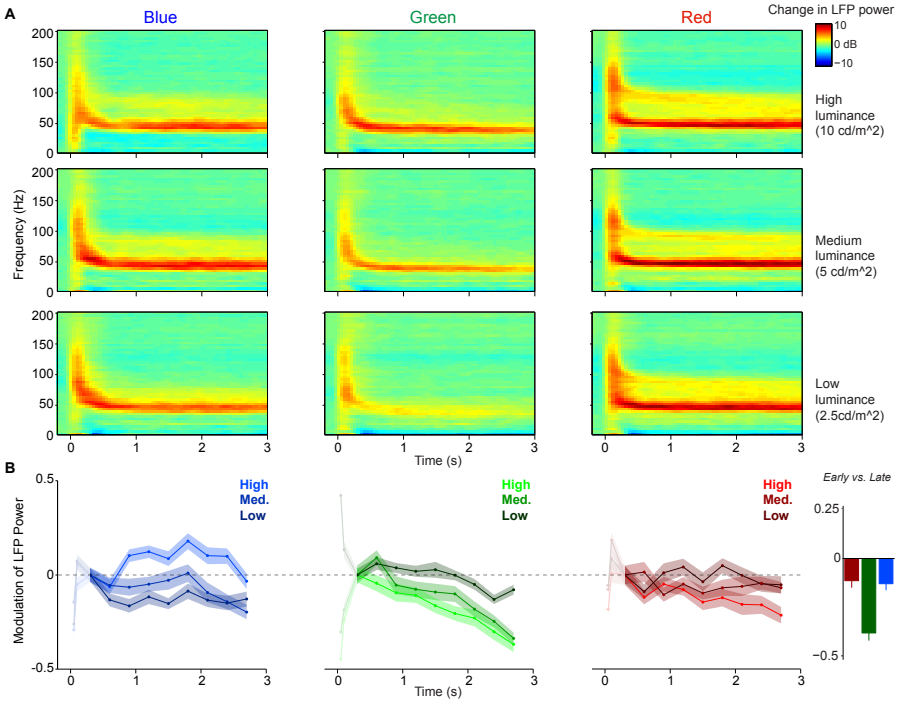

Figure 5.6 — Within-trial temporal dynamics of LFP power spectra during viewing of uniform surfaces. (A) Time-frequency representations of log-transformed change in LFP power relative to baseline (dB), using *Dataset 1* (N = 5 sessions, 19.8±0.45 trials per condition in each session). Shown are the three (equiluminant) different color hues at three different luminance levels. (B) Change in LFP power relative to 0.3-0.6 s period (third time point), separately for different color hues and luminance levels, shown as fold-change modulation index (see Chapter 4 Methods). The first two points correspond to 0.05-0.35 and 0.1-0.4 s and are whitened out because these points are strongly influenced by initial firing transient and associated bleed-in of spiking activity at high frequencies. Right panel shows the modulation of LFP power in early (0.3-0.6 s) versus late (2.7-3.0 s) period, averaged over three luminance levels. The decrease in gamma peak amplitude over time was significantly larger for green than blue and red surfaces (peak amplitude estimated as described in Methods, main effect across luminance levels: $P < 0.05$, bootstrap test), and did not differ between blue and red conditions. This also held true for the modulation of LFP power for the highest luminance condition only ($P < 0.05$, bootstrap test).

5.2.4 Dependence on full-screen background hue

One potential source of adaptation, other than the surface stimulus of a given trial, is the color composition of the continuously presented background. In the experiments described above, all surface stimuli were displayed on a gray full-screen background (FSB). Gamma-band responses to achromatic and chromatic surface stimuli may have been affected by the use of this gray FSB, given that the FSB itself may induce adaptation at many levels of the nervous system. We therefore asked how gamma-band responses to surface stimuli depend on the color of the FSB. To answer this question, we performed experiments in which we used different FSBs in separate, adjacent sessions (gray, white, black, blue, green, yellow and red, Figure 5.7A and Figures 5.11 and 5.8; see Chapter 4 Methods). The FSB was continuously presented during the entire session, i.e. remained on both during the pre-stimulus period, post-stimulus period and the period during which the surface stimuli were displayed (Figure 5.7A). In

Figure 5.7, we analyze LFP responses to the presentation of chromatic surface stimuli of different hues presented at the maximum possible luminance level (see Figure 5.8 for equiluminant red, green and blue as well as achromatic surface stimuli).

We first examined how the responses to surface stimuli with specific hues (e.g. green) were altered by using an FSB with the same hue (e.g. green), comparing them to the sessions with a gray FSB (Figure 5.7B). Surprisingly, we found that when the FSB had the same hue as the surface stimulus, there was a nearly complete abolishment of gamma oscillations for blue, green and yellow stimuli (Figure 5.7B and Figure 5.8). This was observed both when the surface stimulus had the same luminance as the background (Figure 5.7) and when the surface stimulus had a lower luminance (Figure 5.8). Interestingly, red stimuli could still induce detectable gamma oscillations when presented on a red FSB, although the gamma amplitude was strongly reduced compared to the gray FSB condition (Figure 5.7B and Figure 5.8).

The reduction in gamma-band oscillations for the same-hue FSB condition may have been an effect of stimulus size, because the background is effectively a very large surface. Alternatively, it may have been an effect of stimulus history. To investigate these possibilities, we analyzed the post-stimulus period immediately following the offset of a gray surface stimulus that was displayed on a colored FSB. We found that the reappearance of the FSB after the offset of the colored surface induced prominent gamma-band oscillations (see supplemental Figure 5.12).

This indicates that the decrease in gamma-band oscillations with the same hue FSB condition was not due to the large size of the background color stimulation, but that it was due to the continuous presence of the same-hue background. This is also consistent with a previous report showing strong gamma with full-screen color stimuli that change color across trials (Shirhatti and Ray, 2018), similar results from our own recordings (data not shown) and with the positive relation between gamma and stimulus size shown in Figure 4.5C.

Next, we considered interactions between distinct hues. We wondered whether gamma oscillations can not only be reduced by same-hue FSBs, but also enhanced by FSB hues that are different from the stimulus hue, in particular when FSB and stimulus assume opponent colors. The organization of color vision around color-opponency axes, namely the red-green and the blue-yellow axes, is a key principle found both at the neurophysiological and psychophysical level (Wachtler et al., 2003; Livingstone and Hubel, 1984; Solomon and Lennie, 2007; Tailby et al., 2008a). These color opponencies are thought to result from the computation of differences among signals deriving from L and M cones (red-green), and S cones versus L and M cones (blue-yellow). We found that for all surface hues, gamma oscillations were amplified when stimuli and FSBs were of opponent color hues (Figure 5.7C). This suggests that gamma oscillations are dependent on opponency signals along the red-green and the blue-yellow axes (Figure 5.7C). Given the strong dependence of gamma-band oscillations on the FSB, we asked whether the use of a gray FSB may have induced differences in gamma-band amplitude among distinct hues. To examine this possibility, we used a black FSB, which should induce minimal adaptation for all cones. Quite surprisingly, the difference among red, green and blue hues that we had observed with a gray FSB could not be replicated when we presented the stimuli on a black FSB (Figure 5.7D). Compared to the gray FSB condition, gamma-band amplitudes were significantly lower for red and blue surface stimuli and significantly higher for yellow and green surface stimuli (Figure 5.7D). As a consequence, for the black FSB condition, gamma-band power was no longer highest in response to red stimuli, but showed a different dependence on hue (Figure 5.7D; Figure 5.8). Specifically,

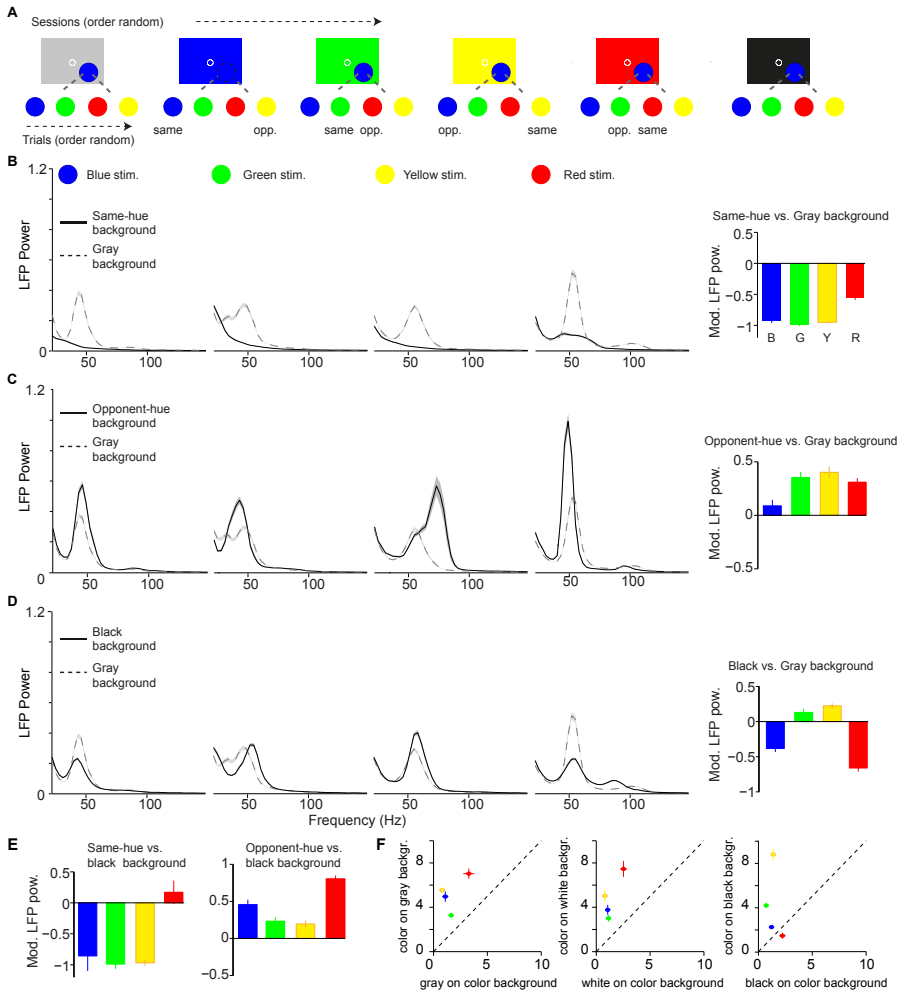

Figure 5.7 — Dependence of LFP power spectra on background stimulus. (A) Illustration of paradigm (*Dataset 5*, see Chapter 4 Methods; n = 25 sessions, 18.64±0.11 trials per condition and session). In a given session, a fixed background stimulus was used, and a set of chromatic and achromatic surfaces (6 or 8 dva) were presented in separate trials (see Chapter 4 Methods). (B) Average LFP power spectra for the different color conditions during gray background versus same-hue background sessions. We show analyses for blue, green, yellow, and red surfaces, presented at maximum possible luminance. Right: modulation index of LFP gamma-amplitudes (see Chapter 4 Methods). Main effect: $P < 0.05$, bootstrap test. R versus G, B or Y, and G versus B: $P < 0.05$, bootstrap test. (C) As (B) for comparison of opponent color background and black background condition. Main effect: $P < 0.05$. B versus R, G or Y, $P < 0.05$. (D) Black background versus gray background. Main effect not significant. All color differences significant ($P < 0.05$, bootstrap test). (E) Modulation of gamma-band amplitude for same-hue vs black background condition (left), as well as opponent-hue vs black background condition (right). Left: $P < 0.05$: R versus G, B or Y; B versus G or Y. Right: $P < 0.05$: all combinations except G versus Y. (F) Comparison of gamma-band responses between chromatic surfaces shown on achromatic background and achromatic surfaces shown on the same respective chromatic background (using the data shown in Figure 5.8).

gamma-band power was higher for green and yellow than red and blue surface stimuli (Figure 5.7D). The resulting pattern could not be explained by luminance contrast differences, because contrast increased for all hues on the black compared to the gray FSB, whereas gamma increased for some hues and decreased for others (Figure 5.7D).

In Figure 5.7B-C, we compared stimulus responses in the same-hue and opponent-hue background conditions with stimulus responses in the gray background condition. However, because of the evidence that the gray FSB may not have affected all stimulus hues equally, we also directly compared the same-hue and opponent-hue FSB conditions with the black FSB condition. This analysis revealed a marked difference between red and the other hues (Figure 5.7E). First, when an FSB of the same hue as the stimulus was compared to a black FSB, gamma was almost abolished for blue, green and yellow, but not for red stimuli. Second, when an FSB of the opponent hue was compared to a black FSB, gamma was enhanced for all colors, but particularly strongly for red (Figure 5.7E). The full matrix of gamma responses for different FSB conditions in Figure 5.8 shows that for all non-red chromatic FSBs, gamma oscillations were strongly amplified for red surface stimuli.

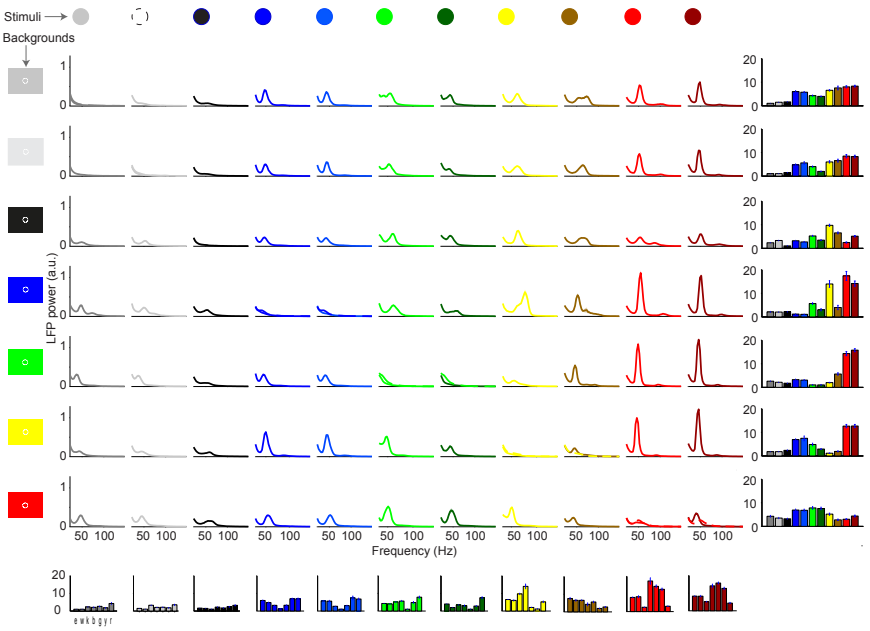

Figure 5.8 — Dependence of gamma LFP power on the combination surface hue and background stimulus. We show here all the condition combinations for *Dataset 5*. Different rows correspond to different stimulus background conditions. The color of the background is indicated by the background stimulus shown on the left. The second row corresponds to a white background. Different columns correspond to different stimulus hue conditions, which are indicated by the color of the lines in each graph. Each graph depicts the average LFP power spectrum, using the same estimation parameters as in Figure 5.7. Bar graphs on the bottom show the gamma peak amplitude (fold-change) for the different backgrounds, separate for each surface hue. Bar graphs on the right show the gamma peak amplitude (fold-change) for the different surface hues, separate for each background condition.

We also analyzed the gamma-responses to achromatic stimuli on colored backgrounds, and asked in particular whether responses of achromatic stimuli on colored

surfaces were as strong as responses of colored stimuli on achromatic surfaces (Figure 5.7F). Achromatic responses on colored backgrounds were substantially weaker than the reverse (see Section 5.2.2 Controls for luminance- and cone-contrast). These data demonstrate that gamma oscillations depended strongly on the FSB, in a way that followed the color opponencies. Furthermore, a commonly used "default" of the display, namely gray, introduced adaptation effects that were color-specific.

5.2.5 A quantitive model relating hue dependence of gamma-band oscillations to adaptation

To explain how gamma-band responses to surface stimuli depend on the FSB, we constructed a quantitative model by estimating the degree to which each FSB differentially adapts the S-, M- and L-cone pathways. Note that this model is agnostic to the neuronal locus at which adaptation of the distinct cone pathways occurs, e.g. it might occur in the retina, LGN or V1.

We hypothesize that gamma-band oscillations for colored surface stimuli are mediated by the activation of cells that are to some degree sensitive to color opponency, in a large spatial region by the same color input. The combination of bottom-up drive at each point of the surface and strong surround modulation may then lead to gamma oscillations. The result that gamma oscillations are particularly strong in the opponent-hue background condition (Figure 5.7) further suggests that when this circuit is more strongly activated (leading to stronger input drive as well as stronger surround modulation), gamma oscillations increase.

Following this reasoning, we further hypothesized that the dependence of gamma-band oscillations on the FSB can be explained by adaptation of specific cone pathways (see Discussion for further argumentation). As an example, a green FSB should lead to stronger adaptation of the M-cone compared to the L-cone pathway. This should increase the degree to which single-opponent cells with L+/M- color-opponencies are activated by red surface stimuli, which may in turn increase the amplitude of gamma-band oscillations.

To capture this intuition in a quantitative manner, we constructed a model in which we aimed to predict the difference in gamma-band amplitudes between red and green surface stimuli (for blue and yellow surface stimuli see further below). The variable to be predicted was the red/green gamma-ratio, defined as $\gamma_{ratio} = \log(\gamma_{red}/\gamma_{green})$, where γ_{red} and γ_{green} are the respective gamma-amplitudes for red and green surface stimuli. This γ_{ratio} was computed separately for all the different FSBs. We estimated the degree to which each FSB adapts the M- and L-cones, using the known response curves of the three cones as a function of wavelength from macaque monkeys (Hárosi (1987), see Chapter 4 Methods, Figure 5.9A).

We then measured the physical wavelength spectrum for each FSB as realized on our monitor. We multiplied the FSB spectra of the different color primaries with the response functions of each cone and summed over wavelengths. This yielded for each FSB stimulus two parameter values, M_{adapt} and L_{adapt}. We then fitted a multiple regression model predicting γ_{ratio} from M_{adapt} and L_{adapt} plus a constant regression intercept (Figure 5.9B). For this model, we used response data for both green and red surfaces presented at maximum possible luminance, as well as equiluminant red and green surfaces, across the different FSBs.

The regression analysis reveals that γ_{ratio} can be highly accurately predicted by the way in which each FSB adapts the L and the M cones (Figure 5.9B; R^2=0.91, $P < 0.05$, F-Test). The regression coefficients for M_{adapt} and L_{adapt} were positive and

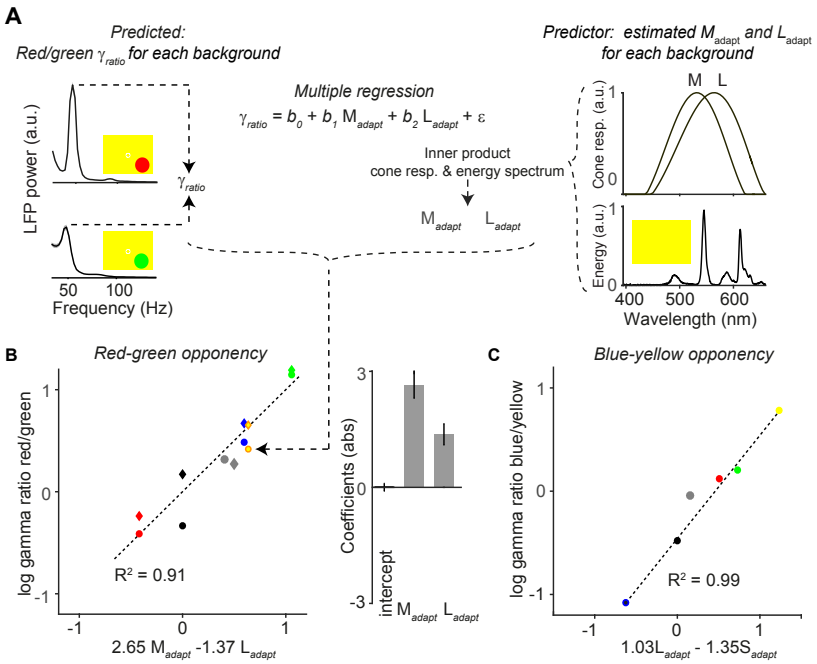

Figure 5.9 — Quantitative model for dependence of gamma-band amplitude on background stimulus. (A) With a multiple regression model, we predicted the ratio of gamma amplitude for red over green surface stimuli. Shown are the LFP gamma power spectra for red (top) and green (bottom) surface stimuli, with a yellow FSB. Normalized cone responses for macaque monkeys are shown on the right, constructed by fitting polynomials to bleaching difference spectra data (Hárosi, 1987). For each full-screen background (FSB) we estimated the extent to which it adapted the three cones (M_{adapt}, L_{adapt} and S_{adapt}; see Chapter 4 Methods), by convolving the spectral energy of the FSB with the normalized cone response curves. For illustration purposes we only show the M- and L-cone response curve. (B) The hues of the data points correspond to the FSB, and the shape indicates whether the red and green surfaces were presented at maximum possible luminance intensity (circle), or whether they were presented at equiluminant intensities (diamond). As predictor variables we used M_{adapt} and L_{adapt} for the different FSBs. The dependent variable was the red/green gamma-ratio, γ_{ratio}. The coefficients indicate the relative influence of the two adaptation parameters on the γ_{ratio}, and the sign indicates whether adaptation of a given cone is increasing or decreasing the red/green gamma-ratio. (C) Similar to (B), but now for blue-yellow. In this case we used S_{adapt} and L_{adapt} as prediction parameters. A model using M_{adapt} in addition did not yield a significant coefficient for M_{adapt}.

negative, respectively. This indicates that adaptation of the M-cone increases γ_{ratio}, whereas adaptation of the L-cone decreases γ_{ratio} (Figure 5.9B). We also found that the γ_{ratio} could not be significantly predicted when using S_{adapt} and L_{adapt} as predictors (P = 0.23), consistent with the idea that the neuronal mechanisms underlying the red-green opponency are dependent on the M versus L cone contrast. The regression intercept of the model (on the γ_{ratio} axis) was not significantly different from zero. This indicates that green and red tend to generate gamma oscillations of similar amplitude when the FSB does not adapt the cones, consistent with the findings shown for the black FSB (Figure 5.7). Strikingly, we found that the M_{adapt} coefficient had an absolute magnitude approximately twice as large as the L_{adapt} coefficient (Figure 5.9B). This suggests that uniform surfaces tend to adapt the M-cone pathway more strongly than the L-cone pathway, or that adaptation of the M-cone pathway has a stronger effect on gamma-band oscillations than adaptation of the L-cone pathway.

The model further explains some non-trivial findings that would have been unexpected if FSBs had affected the M- and L-cone pathway in a similar way: We observed that the yellow FSB strongly amplified γ_{ratio} (Figure 5.9B). Given its wavelength spectrum, the yellow FSB is expected to adapt the L-cones more strongly than the M-cones, which would predict a reduced γ_{ratio}, i.e. red responses being weaker than green responses. By contrast, we found that γ_{ratio} was enhanced. The models explain this by the fact that the stronger L-cone than M-cone activation by the yellow background is more than compensated by the much greater M_{adapt} than L_{adapt} coefficient. Similarly, the γ_{ratio} increased for a gray compared to a black FSBs, even though gray FSBs should in principle adapt the M- and L-cone pathways to a similar degree (Figure 5.7). This was again compensated by the much greater M_{adapt} than L_{adapt} coefficient.

We performed a similar analysis for the yellow-blue (L+M-S) opponency axis, aiming to predict the gamma ratio of blue over yellow (Figure 5.9C). We first fitted a model with the S, L and the M cone parameters, and γ_{ratio} was now defined as $\gamma_{ratio} = \log(\gamma_{blue}/\gamma_{yellow})$. This regression model explained a large degree of variance ($R^2 = 0.99$), with almost equal magnitude of S (negative, -1.23) and L (positive, 1.24) coefficients, but a much smaller and non-significant coefficient for the M cone (-0.29). The finding that the model fit included a highly positive L-cone coefficient and non-significant (and negative) M-cone coefficient seems *prima facie* to contradict the canonical idea that the perceptual blue-yellow opponency axis is mediated by an S versus (L+M) opponency. However, neurophysiological data has shown that the main opponency for LGN cells on the yellow-blue axis is the L versus S cone (Tailby et al., 2008b). We further simplified our model using two predictive parameters, equating the blue axis to the S cone and the yellow axis to the L cone (Figure 5.9C). Again, we found a highly predictive relationship with negative weight for the S cone and a positive weight for the L cone (Figure 5.9C), with only a small and non-significant difference in the magnitude of the adaptation coefficients.

The results were qualitatively highly similar between animals and individually significant (Figure 5.10). These findings indicate that gamma oscillations are mediated not only by opponency signals along the red-green axis, but also along the blue-yellow axes (consistent with the data shown in Figure 5.5). They also indicate differential adaptation of the M- compared to the L-cone.

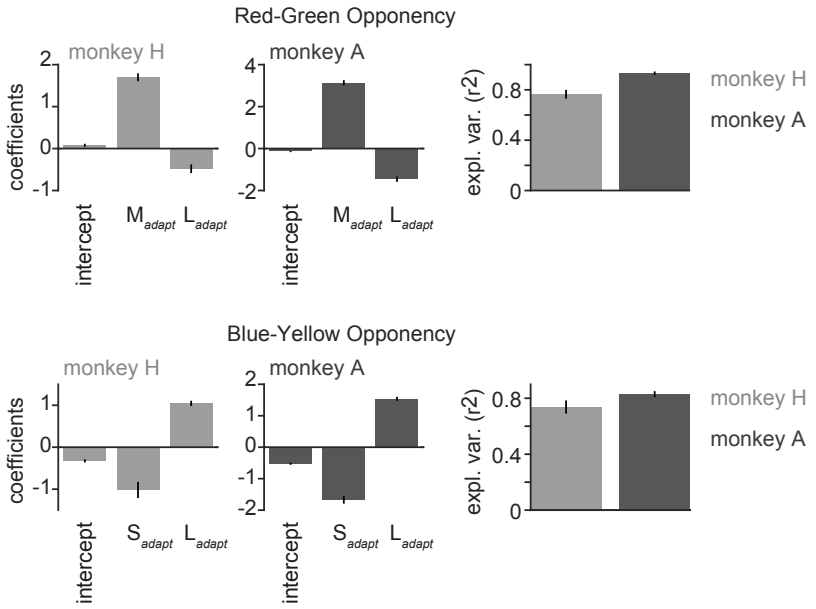

Figure 5.10 — Analysis of 5.9 performed separately for the two monkeys. Top: Red-green opponency. Shown are the regression coefficients of the red-green γ_{ratio} as a function of the adaptation coefficients M_{adapt} and L_{adapt}. Bottom: Same, but now for blue-yellow opponency.

5.3 Discussion

5.3.1 Gamma-band amplitude depends on hue of both stimulus and full-screen background

In the previous chapter, we found a prominent difference in gamma-band synchronization and firing activity between chromatic and achromatic surfaces. Here, we describe prominent differences between hues. A priori, these differences may constitute a constant characteristic of the primate visual system that holds true independently of context. Alternatively, they could emerge from contextualized processing. One major form of such context is the adaptation to the background illumination, in case of the experimental environment the full-screen background (FSB) on which the surfaces were displayed.

Our findings when using a "standard" gray FSB were consistent with previous work using a gray FSB (Shirhatti and Ray, 2018; Rols et al., 2001). Surface stimuli with red (and in our case also pink or brown) hues induced stronger gamma oscillations compared to blueish or greenish colors. Of note, this effect occured in addition to a "main effect" of chromatic versus achromatic stimuli: with a gray FSB, gamma oscillations were reliably induced by all surface stimuli as long as they were chromatic. This is in line with our other findings pointing to contributions of both red-green, but also blue-yellow opponencies to V1 gamma responses (Figure 5.5, Figures 5.7 and 5.9). The study by Shirhatti and Ray (2018) did not report any dependency of gamma-band power on blue-yellow cone contrast. Of note, the stimuli used were not designed to isolate such a cone contrast, and the stimuli varied widely in luminance, similar to our experiment shown in Figure 5.1A. This means that some confounding

covariation of the stimuli in luminance or L-M contrast could have masked any effect of the blue-yellow opponency[2]. Analogously, in our current experiments, we did not find an effect of overall cone contrast (or the related measure of hue saturation). In contrast, the study by Shirhatti and Ray (2018) explored a wider range of hue saturation, and found weaker gamma for perceptually faint stimuli. This is in line with our incidental finding of Figure 5.5A, when attempting to equate stimuli in the red-green and blue-yellow contrast direction simultaneously. This yielded perceptually faint stimuli for the blue-yellow direction, that indeed resulted in very weak gamma responses. In summary, current evidence points to a role for both red-green and blue-yellow contrast, as well as the absolute amount of such contrast for gamma response strength.

The major difference between chromatic and achromatic surfaces and the involvement of color opponencies is also supported by our experiments with changing FSB contexts. In contrast, the particular differences among the specific hues in gamma-band amplitude were highly dependent on the FSB context. Of particular relevance, with a black FSB, and presumably the least adaptation by the FSB through background illumination, red stimuli no longer induced the strongest responses. Gamma to yellow and green increased compared to the gray FSB condition, and gamma to red and blue decreased, despite an increase of luminance and cone-contrast. The pattern of differences among hues elicited for different FSBs was captured well with a quantitative model (Figure 5.9) that allowed differential adaptation of the different cone pathways. The resulting fit suggests that the M-cone pathway adapts more strongly or more quickly than the L- and S-cone pathways. This applied even in cases when the context was perceptually "neutral" (gray FSB).

5.3.2 Is red special?

We infer that background hues during natural vision play a similar adapting role as the FSB did in our experiments. During natural vision, the dominant background hues can vary, and are likely not gray most of the time. In their natural environment, primates typically experience principally green and yellowish stimuli like leaves and trees (Mizokami et al., 2003).

Any claims about a 'special' hue, like strong responses to red, would need to be found robustly across different contexts. In Figures 5.7, 5.8 and 5.9, red shows strong gamma responses in a variety of contexts - as long as some adaptation is involved, i.e. not with a black background. Red stimuli induced the least adaptation, and profited the most from an opponent background (Figure 5.7E). If red is indeed special, how is this achieved? A possible account may be that uniform surfaces induce stronger or faster adaptation of the M-cone than L- and S-cone pathways. This could occur within V1 or in any upstream input, including the retina itself.

The "differential adaptation hypothesis" suggests that if red is special, so will be green due to the red-green opponency. Consistent with this interpretation, gamma-band amplitude decreased more rapidly for green than red or blue surfaces on a time-scale of seconds (Figure 5.6). Additionally, other specific response features of green surfaces are supporting this hypothesis. Namely, a) the luminance dependence was stronger for green stimuli, b) the peak frequency was lower and c) firing rate transients were weaker, but there was no suppression below baseline (Figure 5.1B). This could

[2]In theory, such uncontrolled stimuli can also suggest effects that are the result of underlying confounds.

be interpreted as weaker drive for green stimuli in a gray background condition due to stronger M-cone adaptation. Such a situation would require higher luminance to overcome this adaptation. Similarly, lower peak frequencies are associated with lower stimulus drive for other types of stimuli (Ray and Maunsell, 2010; Hadjipapas et al., 2015; Roberts et al., 2013; Jia et al., 2013b). In the same vein, a weaker transient response could be interpreted as weaker drive, and less suppression as less redundant signal that requires surround interactions involving gamma-band synchronization.

Taken together, these results suggest a mechanism that amplifies the difference between reddish and greenish hues. This has at least two implications that stand out. First, there are considerations for experimental design. Surprisingly, a gray FSB may differentially change neuronal responses to surfaces of different color hues.[3] More generally, a particular design or sampling of hues and backgrounds may greatly affect the results. Hue can interact nonlinearly with luminance (Horwitz et al., 2004). Contrary to a feature space like orientation, it may not be clear how to sample the spectrum of hues evenly: what constitutes a balanced set of conditions? For example, due to the composition of cone spectral sensitivities in old world primates, a larger part of the visible light spectrum will be perceived as green than yellow. Therefore, a simple linear sampling of the light spectrum may induce specific adaptation effects (by cross-adaptation across trials). Another sampling space is suggested by the known color opponencies in the retina and LGN, as implemented in the DKL space (e.g. Figure 5.5). However, this is not identical to the space of perceptual color differences (Wachtler et al., 2003). Given these complexities, experimental designs that purposefully vary the context and/or hue composition, as in Figures 5.1, 5.5 and 5.7, instead of keeping it constant, can provide stronger insights.

Secondly, differential adaptation may affect color perception. Interestingly, at the behavioral level, it has been shown that perceptual after-effects emerge more rapidly after viewing green than viewing red stimuli (Werner et al., 2000). What are the ultimate underlying reasons for the amplification of red-green differences that we hypothesize? Speculatively, this amplification may reflect an ethological need. Given that the visual environment of primates is dominated by green and yellowish stimuli (Mizokami et al., 2003), detection of food such as fruit with relatively high energy in red hues may be an important behavioral task for many primates (Melin et al., 2017). Trichromacy provides behavioral advantages for such detection (Melin et al., 2017), which may be supported by differential adaptation.

5.3.3 Possible limitations and open questions

Given the strong ethological relevance of some hues, is it possible that stronger responses to these hues are a result of stronger salience and therefore attention toward these stimuli? Several lines of evidence speak against this interpretation. First, in Chapter 4, we observed weak gamma responses for sites with RFs covering a salient mismatch position (Figure 4.7). Second, brighter and thereby much more salient hues of red could induce weaker gamma responses than dark red hues (Figure 5.2), pointing toward more stimulus-related mechanisms. A similar argument could be made based on the experiments with changing FSB context, where some highly salient conditions like a black stimulus on a red background nevertheless induce relatively weak gamma-band synchrony.

[3] This surprise is a psychologically interesting phenomenon. A gray FSB hardly gets noticed as a design property, possibly due to its common use and apparent "neutrality", possibly also because such stimuli do not generate strong responses in our visual system.

The present study develops a new "differential adaptation" hypothesis based on a quantitative model of the data. This should have perceptual consequences, and although there is some evidence for such a differential adapting effect on perception (Werner et al., 2000), further studies are required to directly relate behavioral consequences to the neuronal phenomena described here. Further, there could be multiple sources of this adaptation, from the retina to LGN to V1, which require further investigation. Likewise, it is unclear how or if this finding in V1 will transfer throughout the rest of the color system. Given that there seems to be substantial transformation away from red-green dominance from LGN to V1 (De Valois et al., 2000; Wachtler et al., 2003), areas higher up in the visual hierarchy might compensate for the effects in V1. On the other hand, in one of the few studies of this type, it has been shown that adaptation effects in V1 apparently do transfer and affect a downstream visual area (area MT, Patterson et al. 2014). Furthermore, such a compensation higher up in the hierarchy should then likely lead to compensation on the behavioral or perceptual level.

A clear limitation of this chapter is the exclusive focus on surface stimuli. Similar surfaces were used by Wachtler et al. (2003) and Zweig et al. (2015). They are not only of interest from the perspective of their predictability (see Chapter 4), but are also the kind of surface typically used to judge the color of a natural object (Wachtler et al., 2003). Such surface stimuli by definition contain no or very little high spatial frequency (SF) information (i.e. local structure). As discussed in the previous chapter, achromatic stimuli with local, predictable structure, such as a black bar, do induce substantial gamma-band synchronization. What about chromatic, structured stimuli? The data from the chapter 2 show that natural images with color and local structure do induce gamma-band synchronization. On the other hand, using magnetoencephalography (MEG) in humans, Adjamian et al. (2008) showed that gamma-band power was weak or absent for gratings defined purely by their red-green color contrast, i.e. without any luminance contrast. One possible answer to this apparent conundrum is that color is processed in a low-pass fashion, i.e. there is strongest sensitivity to color modulation over a large area of the visual field (Shapley, 2019). Additionally, strictly equiluminant stimulation may be an unnatural condition resulting in weak responses, which reduces detectability of gamma power in MEG, where this detection requires careful stimulus design (Hoogenboom et al., 2006). Further research is needed to establish the spatial-frequency dependence of the hue effects described here.

Finally, the findings described in the last two chapters rely on rather artificial, controlled stimuli under fixation conditions (see also the General Discussion 6). However, in this particular case, these findings are supported by the strong red responses for natural images in Chapter 2. Furthermore, Brunet et al. (2015) studied gamma responses under free-viewing conditions. Although the chromatic/achromatic comparison could not be made in a controlled way with these stimuli, the occurence of clear gamma responses to natural, colored images under free viewing speaks to the relevance of the present findings for natural vision. Remarkably, the strongest gamma observed during free-viewing was for a large orange, an image that has a strong red component with high spatial predictability.

5.4 Chapter summary

This chapter demonstrated that V1 gamma-band responses are hue-dependent and modulated by context in the form of an adapting background.

1. Differences in gamma-band responses among hues exist across a range of luminance differences, and there are strong gamma responses both for stimuli that have red-green (L-M) cone contrast and those that do not (Figures 5.1, 5.5).
2. Red stimuli tend to generate strong responses. This also holds true when cone contrasts are equal to green stimuli, and across a variety of background color contexts.
3. The pattern of hue dependence on the background can be captured by a model that suggests that both color opponent axes (red-green and blue-yellow) are required to predict the pattern of gamma responses.
4. Crucially, the model also suggests stronger adaptation for M- than L-cones, which would generate (as observed) strong red responses under many, but not all, viewing conditions.

5.5 Supplemental for Chapters 4 and 5

(see following pages)

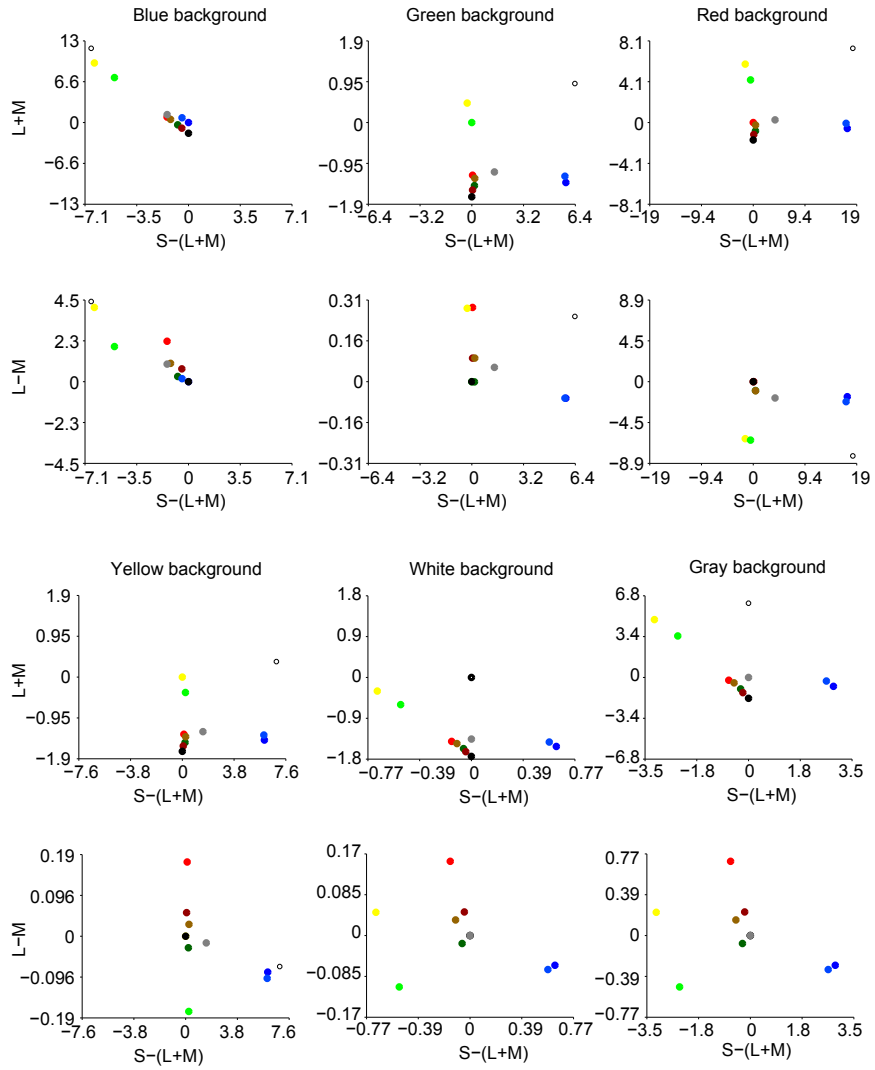

Figure 5.11 — DKL-space representation for Figure 5.7 and Figure 5.8

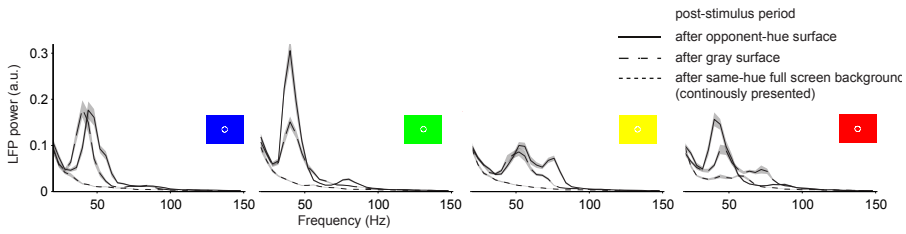

Figure 5.12 — LFP power spectra for the post-stimulus period for the four chromatic background hues (3.5-3.8 s, excluding the initial transient response after stimulus offset at 3.3 s). Clear gamma-band responses were observed for full-screen surfaces after both gray and opponent-hue surface presentation in the stimulus period.

Dataset	Figure	Hue	Lum	CIEx	CIEy
all	all	Background Gray	27.5	347	357
1, 2, 4, 5	1, 2, 4, 6, 7	White	127	343	372
1, 2, 4, 5	1, 2, 4, 6, 7	Black	0.55	315	358
1, 2, 3, 5	all	Red	9.69	624	340
1, 2, 3, 5	all	Green	10	272	555
1, 2, 3, 5	all	Blue	10.7	161	125
1	4, 5	Red	4.77	608	311
1	4, 5	Green	4.41	273	532
1	4, 5	Blue	4.97	163	121
1	4, 5	Red	2.54	574	341
1	4, 5	Blue	2.83	287	513
1	4, 5	Green	2.54	169	129
4	4	Purple	1.35	207	129
4	4	Purple	13.4	206	119
4	4	Purple	12.6	198	114
4	4	Blue	9.42	153	91
4	4	Blue	14.8	162	135
4	4	Cyan	49.4	218	356
4	4	Green-Blue	85.1	277	573
4	4	Green	86.2	281	589
4	4	Green-Yellow	89.8	293	581
4	4	Yellow	110	430	486
4	4	Orange	58.5	520	417
4	4	Red	32.3	641	343
4	4	Red	14.1	631	341
4	4	Red	4.01	593	338
4	4	Red	1.12	494	349
4	4	Brown	28	457	453
4	4	Brown	20.1	510	417
4	4	Purple	13.3	411	218
4	4	Purple	28.2	406	215
5	6, 7	Red	32.3	641	343
5	6, 7	Green	86.2	281	590
5	6, 7	Blue	9.42	153	92
5	6, 7	Yellow	115	425	493

Table 5.1 — Luminances (cd/m^2) and CIE values (*1000) used in this study. Luminances and CIE values were measured with a Konica Minolta CS-100A chromameter, CIE values refer to the 1931 2 degree standard observer. Standard black, white and gray used across datasets are in rows 1-3.

Chapter 6

General Discussion

6.1 Acknowledgements for Chapter 6

A part of this chapter (sections 6.2-6.5) is based on the discussion in Peter et al. (2019), though these sections are more elaborate and integrate information from the previous chapters, and involve solely my own wording, similar to the introductions and discussions of Chapters 4 and 5. The discussion in Peter et al. (2019) was originally written by myself and Martin Vinck, involving extensive discussions. This means that isolating contributions of ideas is not an easy task. I have indicated all ideas that originate from Vinck and Bosman (2016) through appropriate citations, and wish to acknowledge the origin of the notion that gamma-band responses may be poor in a state of inference to Martin Vinck. Some of my own contributions in these sections include the contrasting of evidence from the perspective of predictability versus predictive coding, the notion of gamma-band activity as the reflection of a high-SNR coding regime, the extension of predictability along the temporal domain and the visual hierarchy, the discussion of network mechanisms of gamma-band activity, and pointers to new experiments.

Introductory remarks

A distinguishing feature of V1 activity, induced by many stimulus conditions, is synchronization of neuronal activity in the gamma-frequency band (Vinck and Bosman, 2016; Fries, 2009; Gray et al., 1989; Ray and Maunsell, 2010; Gieselmann and Thiele, 2008). In Chapter 4, the strength of gamma-band power was related to a predictable center-surround relationship, or spatial context. In Chapter 5, prolonged stimulation in time resulted in strong reductions in gamma-band responses. In contrast, repeated stimulation with an identical stimulus could also strongly increase gamma-band responses, as seen in Chapters 2 and 3.

In the following sections, the role of both spatial predictability and drive in the generation of gamma-band oscillations will be considered first. Thereafter, the role of temporal and spatial context will be compared. Underlying generative mechanisms and the functional role of gamma-band oscillations will be considered in the light of the new findings. The chapter concludes with some methodological considerations, the role of technological advance and an outlook.

6.2 Gamma synchronization, spatial predictability and drive

"The amount of information being fed to the central nervous system must be enormous. [...] Storage and utilization of this enormous sensory inflow would be made easier if the redundancy of the incoming messages was reduced. [...] Physiological mechanisms [...] appear to have arisen by evolutionary adaptation of the organism to types of redundancy which are always present in the environment of the species. Much of the sensory input is not shared by all individuals of a species [...] so a device for 'learning' to reduce redundancy is required." Barlow (1959)

Gamma oscillations have been related to sensory predictions and prediction errors based on at least two opposing hypotheses (Chalk et al., 2016; Jadi and Sejnowski, 2014; Vinck and Bosman, 2016; Bastos et al., 2012; Arnal and Giraud, 2012): 1) Bastos et al. (2012) and Arnal and Giraud (2012) posited that gamma-band activity mediates the encoding and transmission of prediction error signals from lower to higher areas. In line with this notion, feedforward influences among visual areas are strong in the gamma band (Bastos et al., 2015) and gamma-band activity has been reported for mismatches of auditory and visual information (Arnal et al., 2011). 2) In contrast, Vinck and Bosman (2016) postulated that the emergence of gamma oscillations depends on predictable relationships between the RF and the surround. Using the concepts of RF and surround, the "predictability hypothesis" thereby links gamma oscillations to theories of efficient and (some variants of) predictive coding, which suggest that feature-specific surround modulation serves to generate a sparse but informative response to the visual input (Rao and Ballard, 1999; Barlow, 1959; Coen-Cagli et al., 2015; Schwartz and Simoncelli, 2001).

In Chapter 4, it was shown that a chromatic mismatch at the center of a uniform surface strongly reduced the amplitude of gamma oscillations for sites with RFs at the mismatch location, yet led to a strong increase in their firing activity. In contrast, gamma oscillations remained strong at nearby sites with RFs in the uniform regions of the surrounding surface (Chapter 4, Figure 4.7). In other words, when the stimulus input at the center of the stimulus (e.g. a red blob) was not predicted from the surround (a green surface), gamma oscillations were weak at the center, but firing activity was strong. This is in line with the "predictability" hypothesis, rather than the notion that gamma-band activity carries prediction errors. Similarly, stronger gamma-band responses with increasing stimulus size (see Chapter 4) are interpretable from the perspective of center-surround interactions, whereas it is less clear why a larger stimulus should be less predictable and therefore produce prediction errors. The seminal studies of Gray et al. (1989) showed that a bar stimulating two distant receptive fields synchronized neuronal responses, whereas two smaller, aligned bars targeted on the two individual RFs did not. This result is also more in line with the idea of predictable relationships, rather than prediction error, generating gamma-band synchronization. As a last example, Jia et al. (2013b) demonstrated that the superposition of dynamic noise on a grating stimulus can reduce gamma-band responses. It therefore seems that "predictability" of center-surround relationships may constitute a necessary ingredient for gamma-band synchrony, both locally and between distant RF locations.

Recently, Auksztulewicz and Friston (2016) interpreted the gamma-band increases with repetition (Chapters 2 and 3, Brunet et al. 2014) as reflecting "precision" of prediction errors in terms of "predictive coding" (Friston, 2005). Since gamma-band increases were related to increased synchronization of putative interneurons, whereas most putative excitatory cells decreased their locking to the gamma-band, they rea-

soned that gamma-band increases reflected gain control. Increased gain control is related to increased precision. One should note that "precision" reflects the confidence in a prediction, which seems conceptually close to a notion of "predictability", but focused on the temporal domain. This also suggests that an overall increase in gamma-band activity, visible in the LFP or noninvasive recordings, can reflect both increases and decreases in prediction error in this "predictive coding" framework, and cannot be interpreted without knowledge of the behavior of different cell types. In any case, this indicates that gamma-band repetition effects can be interpreted along the lines of temporal predictability (see also next section).

Next to "predictability", a second central element that has been previously related to the strength of gamma-band responses is input drive. The dependence of gamma-synchronization on input drive is a feature of many computational models and is illustrated by the dependence of gamma-band power and peak frequency on luminance contrast (Börgers and Kopell, 2008; Roberts et al., 2013; Lowet et al., 2015; Hadjipapas et al., 2015; Ray and Maunsell, 2010; Jia et al., 2013b; Henrie and Shapley, 2005; Tiesinga and Sejnowski, 2009; Whittington et al., 2011). Two factors which had strong influences on gamma-band power in the previous two chapters could be interpreted in terms of drive. Namely, 1) whether the surface was chromatic or achromatic, and 2) the full-screen background on which the surface was displayed. As discussed in Chapter 4 (section 4.5.1), differences in gamma-band response strength for chromatic versus achromatic stimuli may be related to different processing strategies. However, these stimuli likely also differ in overall input drive. For example, Chen et al. (2007) described a stronger source of drive for chromatic stimuli in LFP signals along the ventral stream, including V1.

Responses to chromatic stimuli could be weakened considerably by prolonged adaptation with a full-screen background of the same hue (Figure 5.7). Gamma oscillations could also be strongly amplified, namely by using a full-screen background of an opponent hue (Figure 5.7). The quantitative model (Figure 5.9) captures these findings with respectively a decrease (same-hue background) or increase (for opponent-hue background) in the activity of a cone-specific pathway, i.e. changes in drive. While previous modeling and experimental work has considered the role of drive, it was not part of the original proposal of the "predictability" hypothesis, and constitutes a refinement of the original proposal by Vinck and Bosman (2016).

It therefore seems that gamma-band oscillations depend on both spatial "predictability" and sufficient drive. Notably, these two factors have opposite correlations with neuronal firing rates. Taking into consideration both center-surround relationships and the input drive may help to better understand previous reports that gamma-band responses can be both positively and negatively related to firing rate (Jia et al. 2013b; Gieselmann and Thiele 2008; Hadjipapas et al. 2015; Roberts et al. 2013, see Jadi and Sejnowski 2014 and Jia et al. 2013b for examples of related modeling work). Intuitively, sufficient drive will generate enough signal that has the potential to be explained away if predictable. This relates gamma-band activity to a particular stimulus regime and also a particular goal from the perspective of predictive coding. As described in the General Introduction (sections 1.3.2 and 1.4.2), an understanding of the visual world that uses spatial context can entail either inference of ambiguous information, or efficient coding through explaining away of redundancies. These two aspects are related, since both have been suggested to require a generative model of the sensory data (Rao and Ballard, 1999; Friston, 2005; Coen-Cagli et al., 2012). However, inference has been related to facilitation from a predictive surround, whereas efficient responses are related to surround suppression. These two operations, statistical infer-

ence and coding efficiency, could therefore rely on distinct circuit mechanisms.

In a high signal-to-noise regime, the removal of redundant (predictable) information should lead to an efficient code, likely involving inhibitory mechanisms (Jadi and Sejnowski, 2014; Vinck et al., 2013b) accompanied by gamma-synchronization (see Mechanisms section below). If the local signal is weak, inference of the stimulus input at one location using contextual surround information may rely predominantly on excitatory surround influences (De Lange et al., 2018; Rao and Ballard, 1999; Coen-Cagli et al., 2012). An excitatory influence of the surround has been shown in cases where the input into the classical RF provides a weak drive, e.g. collinear facilitation of low-contrast but not high-contrast bars by a matching surround (Kapadia et al., 1995, 1999), and responses to illusory contours (von der Heydt et al., 1984; Grosof et al., 1993). Given the strong role of inhibition in the generation of gamma, this regime may not be conducive for the emergence of gamma. It might be very informative to explicitly measure gamma-band synchrony in stimulus situations like the collinear facilitation experiment to test the notion that gamma is absent in the case of inference. Notably, successful inference entails some surround predictability. It should also be noted that in the case of a large, highly predictable grating stimulus, Michelson contrasts below 10% yield detectable gamma-band responses (e.g. Hadjipapas et al. 2015): the effect of drive appears more graded than the rather qualitative change in gamma-band responses for center-surround relationships and stimulus size.

Taken together, this suggests a predictability-and-drive regime for gamma. Notably, there are species differences in the expression of gamma in the visual system. Could a predictability-and-drive regime be species-specific? First, can there be higher predictability of the same environment for different species? Gamma-band activity is strong in humans, other primates and carnivores, in particular cats (reviewed in Vinck and Bosman 2016). All these species have high visual acuity, and small receptive fields. The visual world is highly redundant, or predictable, at nearby locations (Barlow, 1959; Marr, 1982; Field, 1987). As pointed out by Vinck and Bosman (2016), for species with small RFs and high visual acuity, "predictability" will therefore be a frequent occurence, in particular in early visual areas. Notably, high-acuity species, including humans, also share highly structured horizontal connectivity patterns and long-range connections in V1, which may aid in "predictability"-related computations (Kaschube et al. 2010, see Mechanisms section).

What about differences in drive? Speculatively, it is worth noting that a major determinant of a high signal-to-noise environment in vision is sufficient lighting. Color perception is a hallmark of daylight vision, which generates strong drive to the visual system and strong gamma-band responses (Chapters 4 and 5). Free-viewing of colored natural images can generate strong gamma-band responses (Brunet et al., 2015). This hints to the fascinating possibility that species with high-acuity visual systems that experience high-SNR regimes may adapt either phylogenetically or on shorter timescales to generate synchronized, highly efficient responses. This may possibly even involve specialized cell types (see Mechanism section below). For a predictability-and-drive regime, the drive aspect could therefore stem from diurnal activity. Alternatively, relatively poor drive could be amplified, as in crepuscular species like cats that possess a tapetum lucidum[1].

In summary, gamma-band synchronization shows dependencies on both "predictability" and drive. The combination of these two factors relates gamma to a particular

[1] A tissue in the eyes that lies behind the retina and acts as a reflector. It should be noted that cats studied in the laboratory, like other cats interacting with humans, will be active during the day.

stimulus and coding regime, characterized by surround suppression and efficient coding. Species with high visual acuity will likely experience this regime frequently, such that gamma-band synchronization may be optimized on phylogenetic or developmental timescales.

6.3 Effects of spatial and temporal context on gamma synchronization

"No sensory stimulus is an island unto itself; rather, it can only properly be interpreted in light of the stimuli that surround it in space and time."
Schwartz et al. (2007)

The previous chapters relate to both spatial and temporal contextual effects on gamma-band synchronization. In Chapter 4, it was shown that when the RF stimulus is part of a larger uniform surface, and the spatial context allows a prediction of RF content, gamma oscillations are enhanced. What about temporal predictability? In Chapter 5, in the adaptation paradigm, the RF stimulus was part of a longer uniform stimulation period, and temporally predictable. However, gamma oscillations were strongly reduced. In contrast, the repeated (and therefore to a certain degree predicted) presentation of a stimulus, as another form of temporal context, can increase gamma-band oscillations (Chapters 2 and 3).

Taken together, this suggests that these effects are brought about by different mechanisms. The case of continuous stimulation will be considered first. As discussed in Chapters 2 and 5, prolonged stimulation with an identical stimulus can lead to adaptation (reduction) of the drive to V1. In this case, there is insufficient signal to engage the interaction between excitatory and inhibitory cells that underlies gamma-band oscillations, and there is likely less redundancy in the remaining signal. In contrast, previous studies have measured responses to stimuli with varying temporal predictability that are not adapting the inputs as strongly. In cases where continuous stimulus motion has a large degree of jitter or randomness, V1 gamma-band synchronization tends to be weak. In contrast, it tends to be strong in cases where stimulus motion is predictable (Kayser et al., 2003; Kruse and Eckhorn, 1996; Gray et al., 1989; Besserve et al., 2015a), as is V1-V4 gamma-band coherence (Bosman et al., 2012; Grothe et al., 2012). The stimuli that generated strong gamma-band responses had a predictable spatial structure and a high SNR, and additionally showed some predictable temporal behavior. They may therefore constitute a temporal extension of the drive-and-predictability regime. More generally, these findings are in line with a role also for temporal predictability in the generation of gamma-band responses.

However, previous work using MEG and EEG recordings in humans have suggested that when stimuli were unexpected given the immediate temporal or cross-modal context, MEG/EEG gamma power increases (Todorovic et al., 2011; Arnal and Giraud, 2012). This led to the suggestion that gamma-band activity mediates the transmission of feedforward error signals (Bastos et al., 2012). Several possible explanations for these phenomena can be put forward. Firstly, as discussed in Chapter 2, some form of rapid adaptation and thereby a decrease in drive might occur for a stimulus repetition. This rapid process has opposite effects on gamma-band power compared to the repetition effects with repeated stimulation described by Brunet et al. (2014). Therefore, in some cases, gamma-band increases for expectation violations may reflect rapid adaptation to matches, rather than stronger responses to mismatches

(Kaliukhovich and Vogels, 2014). A second possibility is that expectation violations elicit so-called high-gamma activity. High-gamma activity has been reported for various stimuli and cognitive tasks. It tends to be transient rather than sustained, has been related to learning and plasticity by its NMDA dependence, and may reflect dendritic processes (Bartoli et al., 2019; Leszczynski et al., 2019; Ray and Maunsell, 2011). Broadband high-gamma responses are therefore likely to be observed during cognitive processing, including prediction errors. Thirdly, another interesting possibility is that a top-down feedback error signal may be generated for expectation violations, possibly as a teaching signal to improve responses in lower visual areas. Compared to predictive coding theories, error signals would travel in the other direction in this scenario (Heeger, 2017). Evidence for such error signals originating from higher visual areas has been reported recently for interactions between higher and lower-order areas within IT (Schwiedrzik and Freiwald, 2017; Issa et al., 2018). Speculatively, this feedback could generate drive in superficial layers in the lower area of IT, that may be sufficient to generate a gamma-band response. Areas in IT no longer possess a clear spatial surround, but feedback drive may stimulate the "effective surround" in a higher visual area, i.e. neighboring, connected cells, possibly sharing feature preferences. Interestingly, *temporally* stimulating such a feature-based surround in IT can show similar effects for rates as a *spatial* surround in V1 (Kaliukhovich and Vogels, 2016).

At first glance, spatial and temporal context might appear as separate aspects of vision. However, natural vision contains simultaneous dependencies in space and time, for example for any moving object. Kaliukhovich and Vogels (2016) elegantly showed that concepts developed for surround modulation in early visual cortex could be extended to IT in the temporal domain. Similarly, temporal and spatial context may interact strongly within V1 (Aschner et al., 2018; Schwartz et al., 2017). Recent modeling work show that a normalization of responses by predictable spatial context can be extended to a predictable temporal context. This captures adaptation and surround modulation with a similar operation, and with a similar result of redundancy reduction (Snow et al., 2016). I therefore suggest that a full understanding of of gamma-band responses may require a concept of a surround that applies all along the ventral stream, as well as in space and time. Based on this extension of the "predictability hypothesis", I suggest that using naturalistically moving or transforming stimuli may be a way to generate stronger gamma-band responses in higher visual areas. Possibly, this might overturn the current state of the literature that suggests a decrease of gamma-band response strength with increase in hierarchical position along the ventral stream (reviewed in Vinck and Bosman 2016).

The temporal surround will integrate information on short timescales. Immediate, single repetitions or constant stimulation can easily lead to input fatigue, and are particularly easy cases to predict that do not require additional computation after some point. In contrast, predictable stimulus motion as during natural vision allows for some useful computation of future events, which may be a key operation of the visual system (Hénaff et al., 2019). Since areas higher along the ventral stream can detect higher-order correlations compared to V1 (Yu et al., 2015; Freeman et al., 2013), it is possible that more complex forms of "predictability" drive higher areas in the ventral stream (Barlow, 1959).

On a longer timescale, a predictable center-surround stimulation via repeated presentations might encourage an optimization of responses to these stimuli (Brunet et al. 2014, Chapter 2). With regard to gamma-band responses, increases with stimulus repetition may reflect a short-term learning process in which center-surround interactions

are modified by experience. This could involve strengthening the effective connectivity between cells that have strong responses to the stimulus between center and surround. It would be interesting to test this hypothesis using targeted stimulation of cells in the center or surround, or possibly by the analyses of millisecond timescale interactions using coupled oscillator theory (see Mechanism section below).

To summarize, current evidence is compatible with an extension of the drive-and-predictability regime for gamma-band activity into the temporal domain. Prolonged stimulation may lead to fatigue and reductions in drive, whereas naturalistic motion combines drive and predictability in time. Moreover, repeated stimulation could result in both short-term and long-term learning of spatially predictive relationships.

6.4 Mechanisms of gamma synchronization

"Theorizing at this stage is like skating on thin ice - keep moving or drown."
Hebb (1949), p.xii

The previous chapters and discussion point to some of the principles underlying the stimulus- and contextual dependence of gamma synchronization. Yet, the precise neuronal mechanisms generating V1 gamma synchronization, and in particular its dependence on spatial and temporal context, remain unclear. These mechanisms include their origin, cellular substrates, and network synchronization principles.

Origins. Previous work indicates that in primate and cat V1, gamma-band oscillations are generated cortically, with especially strong generators in the superficial layers of the cortex (Xing et al. 2012; Buffalo et al. 2011; Livingstone 1996; Herculano-Houzel et al. 1999; Bastos et al. 2014, see also General Introduction, Figure 1.4). Both lateral connectivity and top-down feedback is strong in superficial layers (Lund et al., 1993; Barone et al., 2000; Markov et al., 2014). Early reports in anesthetized animals suggested the presence of gamma-band oscillations in the LGN that had a distinct frequency profile from cortical oscillations, pointing to independent cortical generators (Castelo-Branco et al., 1998; Neuenschwander and Singer, 1996). However, a recent study did not detect gamma-band oscillations in the LGN of awake primates (Bastos et al., 2014). In Chapter 4, the dependence of gamma-band synchronization on a predictable center-surround relationship was illustrated, in line with previous findings of the size dependence of gamma-band oscillations and the emergence of gamma-band oscillations for bar stimuli that continue in a predictable manner into the surround (Gieselmann and Thiele 2008; Jia et al. 2011; Chalk et al. 2010; Gray et al. 1989, reviewed in Vinck and Bosman 2016). Taken together, this suggests gamma oscillations emerge in superficial layers, given the integration of bottom-up inputs with contextual information mediated through lateral and top-down feedback.

Cellular underpinnings. Assuming therefore that the origin of gamma-band oscillations is within cortex, we can next consider the role of different cell types. Currently, it is common consensus that the interaction between inhibitory and excitatory neurons likely plays a central role (Jadi and Sejnowski, 2014; Veit et al., 2017; Hasenstaub et al., 2005; Buzsáki and Wang, 2012; Kopell et al., 2000; Whittington et al., 1995; Tiesinga and Sejnowski, 2009; Perrenoud et al., 2016; Cardin et al., 2009; Bartos et al., 2007; Sohal et al., 2009; Womelsdorf et al., 2014; Vinck et al., 2013a). How such interactions could generate oscillations was introduced in the General Introduction, section 1.4.3. Both computational and experimental work has implicated fast-spiking basket cells in the generation of cortical gamma oscillations (Moore et al.,

2010; Cardin et al., 2009; Vinck et al., 2013a; Buzsáki and Wang, 2012; Bush and Sejnowski, 1996; Wang, 2010). Interestingly, it has been shown that surround modulation increases the firing of fast-spiking interneurons in cat visual cortex (Haider et al., 2010). Stimulation of PV interneurons, a type of inhibitory cell that expresses parvalbumin and provides highly effective inhibition directly to the soma of excitatory neurons, can generate gamma-band responses in cortex (Cardin et al., 2009). PV neurons have long been suggested to be central for gamma-band activity, and have been related to gain control due to their perisomatic inhibition.

However, for V1 in particular, the views on which interneuron type may be most crucial for gamma oscillations have been evolving in recent years. Cell-type specific manipulations are increasingly feasible, particularly in rodents. In mice, somatostatin-positive SOM interneurons, targeting dendrites rather than soma of excitatory neurons, were found to better correlate with visually induced gamma-band activity (Moore et al., 2010; Chen et al., 2017; Veit et al., 2017). Not only did driving these cells induce gamma as in PV neurons, suppressing SOM cells reduced gamma-band activity, an effect not seen for PV cells (Veit et al., 2017). This is interesting because SOM interneurons target dendrites, where gamma rhythms could gate computations or plasticity more flexibly than at the soma, by targeting specific inputs. SOM interneurons have long-range horizontal connections and stimulus size-dependent firing properties that are pointing to a role in contextual integration. Since driving both PV and SOM neurons can increase gamma-band responses, inasmuch as a function can be mapped onto a neuron type, gamma-band activity could reflect both contextual integration and gain control (related to "predictability" and drive, respectively). Notably, the oscillations in mice occur at frequencies below 30 Hz, and it is currently unclear which types of stimuli and possibly learning histories result in these responses, since studies with relatively similar stimuli do not always report these oscillations (Veit et al., 2017; Perrenoud et al., 2016; Vinck et al., 2015).

An apparently specialized class of excitatory neuron, the "chattering" (i.e. fast-rhythmic bursting) cell, deserves special consideration for V1 gamma rhythmogenesis. These cells generate gamma-range burst firing when artificially depolarized with a current injection (Gray and McCormick, 1996) and show resonant properties in the gamma-frequency band (Gray and McCormick, 1996; Cardin et al., 2005; Nowak et al., 2003). Such cells were reported to be highly prevalent in V1 in the 1990s and early 2000s, but this research field was largely dormant until recently (Onorato et al., 2019). However, it may be that a large portion of V1 cells are "chattering", also in awake primates (Onorato et al., 2019). These cells show stronger stimulus selectivity (in terms of orientation preference) as well as stronger locking to the gamma rhythm compared to other cell types. Interestingly, this cell type may be specific to cats and primates and appears absent in rodent V1, where gamma oscillations typically have a frequency about half of that in primates. Important future questions include whether "chattering" cells show similar properties for chromatic stimuli, what their surround modulation behavior is, particularly with respect to spatial predictability, and how they might behave for repeated stimuli.

To summarise, various cell types have been associated with the generation of gamma-band activity. These include two types of interneurons whose activity has been related to contextual integration. Recent findings of a very stimulus-informative, gamma-locked excitatory cell type in highly visual carnivores and primates support the proposal that that may reflect a specialization to a predictability-and-drive regime these animals are likely to experience.

Network mechanisms. Regardless of the underlying cell types, it is necessary to

consider gamma-band activity as a network phenomenon. Many computational models can account for the generation of local gamma oscillations (Börgers and Kopell, 2005; Whittington et al., 2011; Börgers and Kopell, 2008; Tiesinga and Sejnowski, 2009). These models capture the general necessity of sufficient input drive, and also drive as one explanatory factor for the stimulus dependence of gamma-band oscillations (Börgers and Kopell, 2008; Lowet et al., 2015). However, a strong drive alone is apparently insufficient to generate narrow-band gamma oscillations, as discussed above. It remains largely unknown what mechanisms are underlying gamma-band synchronization between more distant V1 columns with non-overlapping RFs. Recently, it has been shown for awake primate V1 that a lot of the apparent stochasticity and fluctuations in gamma-band frequency that occur on a millisecond time scale can be captured using the well-established theory of weakly coupled oscillators (TWCO, Lowet, Roberts, *Peter*, Gips and de Weerd, (2017)). TWCO describes synchronization as a dynamic process, where oscillators that show sufficient coupling strength modulate one another's instantaneous frequency. Moment-to-moment interactions between distant V1 sites could be captured using this approach, and a dependence on two main parameters of TWCO, coupling strength and difference in frequency between sites, was established. Since stimulus drive can modify oscillation frequency, for example via luminance contrast (Ray and Maunsell, 2010; Roberts et al., 2013; Hadjipapas et al., 2015; Jia et al., 2013a), desynchronization between distant sites may occur when visual inputs provide a different stimulus drive. Next to frequency tuning, coupling strength (e.g. existing connectivity strength) will determine synchronization. Tangential, excitatory connections linking columns with similar feature preferences (e.g. color or orientation) may play a crucial role in synchronizing neuronal assemblies coding for related features by coupling local oscillators (Vinck and Bosman, 2016; Gray et al., 1989; Korndörfer et al., 2017; Lowet et al., 2017). This suggests that "predictability" on the level of V1 requires horizontal connectivity, or coupling. For example, the experiment by Gray et al. (1989) showed that a bar connecting distant RFs generated gamma-synchronized spiking, whereas – arguably predictable – aligned, small bars in the individual RFs did not.

A stimulus with high spatial predictability may simultaneously activate a large number of excitatory horizontal connections, which could then give rise to gamma-synchronization via the recruitment of local excitatory and inhibitory neurons. In addition, feedback from higher visual areas could be critical, considering that the spatial spread of tangential connections is somewhat limited and covers a smaller surround region than cortical feedback (Angelucci et al., 2017). Currently, the relative contributions of local, horizontal connectivity compared to top-down feedback are unclear.

Importantly, there is evidence that predictable patterns are learned from natural input statistics and reflected in network connectivity strength (Fiser et al., 2010; Berkes et al., 2011). Speculatively, less predictable patterns might tap into less strongly coupled populations, which are unable to synchronize even given similar input strength. Stimulus repetition could affect effective connectivity strength, for example through synaptic depression or facilitation, resulting in stronger gamma-band responses with repetition if the structure of the image can be learned by the circuit. By extension, long-term learning could change structural connectivity in response to stimuli to allow more efficient, cooperative encoding. This predicts the emergence of stronger gamma-band responses for novel stimulus classes over the course of days, if stimuli are sufficiently novel to require a restructuring of cortical responses, and endowed with structure that cortex is able to extract.

Finally, one alternative view to the idea of coupled oscillators is that gamma-band

responses behave like a traveling wave (Besserve et al., 2015b; Muller et al., 2018). This traveling wave may only emerge when there are predictable input relationships over a large region in space, with a locally measured oscillation being a snaphot of such a wave (Muller et al., 2018).

To summarise, the origins, cellular basis and network mechanisms all point to a role of gamma-band synchronization for contextual integration in a predictability-and-drive regime.

6.5 Functions of gamma synchronization

"Cortical computation cannot be understood without the consideration of cortical dynamics, and cortical dynamics cannot be understood without the consideration of cortical computation." Vinck and Bosman (2016).

The study of visual gamma-band oscillations encompasses not only a dissection of the mechanisms as described above, but from its earliest days also ideas about its functional relevance. What are the functional implications of the previous chapters?

Binding by synchrony interpretations. Early theories maintained that gamma synchronization could overcome the so-called "binding problem" (Singer and Gray, 1995; Singer, 1999). This refers to the problem of segmenting images into objects, with different features and parts of the objects grouped together. Given the distributed activity in the visual system in response to a stimulus, it was suggested that the activity can be dynamically grouped together through synchrony according to perceptual Gestalt principles (Milner, 1974; Engel et al., 1992; Singer and Gray, 1995; Singer, 1999; von der Malsburg, 1994; Singer, 2018). However, it was soon found that some stimuli that follow Gestalt principles do not induce gamma-band synchrony, for example between connected sites that exceed the maximal horizontal connectivity in V1 (Roelfsema et al., 2004). Simple stimulus manipulations have strong effects on gamma-band responses (Lima et al., 2010; Jia et al., 2013b, 2011; Hadjipapas et al., 2015; Ray and Maunsell, 2010; Peter et al., 2019), such that neurons responding to different parts of the same object do not necessarily show synchrony. Therefore, binding-by-synchrony cannot be the general solution as originally proposed. However, notably, some operations that may contribute to perceptual grouping and underlie some of the Gestalt principles, such as contour integration (Liang et al., 2017), perceptual filling-in (Zweig et al., 2015; Land, 1959; Wachtler et al., 2003), and figure-ground segregation (Lamme, 1995), have been linked to surround modulation. A link between center-surround predictability and gamma-band synchronization as discussed above may therefore relate gamma-band synchrony to proto-operations for object recognition.

Communication through coherence. Whereas early work focused on the binding problem, some later work switched the emphasis to flexible coordination ("regulation of communication") between neuronal populations, especially between brain areas (Fries, 2005; Salinas and Sejnowski, 2001; Palmigiano et al., 2017; Akam and Kullmann, 2010; Knoblich et al., 2010; Colgin et al., 2009; Jia et al., 2013a). For example, the communication-through-coherence (CTC) hypothesis states that cognitive demands can selectively adjust coherence between neuronal populations to enable or prevent information transmission (Fries, 2005, 2015, 2009). This involves the gating of inputs to downstream areas by the ongoing rhythm in the downstream areas, or coordinated rhythms between areas. Indeed, recent studies have shown that distant

visual areas show selective gamma-band coherence, primarily when they process an attended stimulus (Gregoriou et al., 2009; Bosman et al., 2012; Grothe et al., 2012; Buschman and Miller, 2007). The level of coherence could predict reaction time (Rohenkohl et al., 2018). Furthermore, Jia et al. (2013a) directly measured the influence of spiking activity in V1 on the highly connected, downstream area V2 in the presence of stimulus-induced gamma band oscillations. Impressively, they showed that for a doubling of pairwise synchrony in V1 responses, a V1 action potential was about twice as likely to elicit a spike in V2. However, they also reported that this effect seemed to mainly depend on synchrony within V1, and not a coordination between V1 and V2. Of note, these data were recorded under anesthetized conditions, possibly preventing the top-down coordination suggested by findings in attention paradigms. The current work cannot add directly to this line of research, as interareal coordination was not studied. Rather, the preceding chapters emphasize the strong stimulus dependence, and the role of spatial and temporal context in the generation of gamma-band responses. To the extent that the observed increases in gamma-band power and LFP-MUA locking are reflecting pairwise synchrony, the findings by Jia et al. (2013a) do suggest that the nature of the stimulus and spatial or temporal contextual modulations in gamma-band responses will have strong effects on downstream areas. Any flexible coordination based on rhythmic synchronization will need to accommodate for these effects. One pathway to address these challenges may be to actively manipulate stimulus attributes, in particular center-surround relationships, whilst recording interareally during an attention task.

Recently, Vinck and Bosman (2016) proposed two, largely untested, functional roles of V1 gamma-band activity in the context of efficient and predictive coding. The first relates to interareal coordination. Vinck and Bosman (2016) suggest, in line with previous suggestions and Jia et al. (2013a), that gamma synchronization may be a mechanism to increase the effective synaptic gain of V1 neurons on post-synaptic targets (e.g. V2) - crucially, when a stimulus is efficiently encoded (Vinck and Bosman, 2016). This may ensure reliable transmission of V1 outputs in circumstances when firing is sparse. For example, a large, colored stimulus can suppress neuronal firing below baseline levels (Chapter 2, Chapter 4). Nevertheless, the increase in gamma-synchronization might ensure that stimulus information is still transmitted with high fidelity. In other words, some stimuli may elicit a synchronization-based, and other stimuli a rate-based encoding strategy in V1. An effective strengthening of gain from lower to higher areas would also accommodate the finding that gamma-band influences tend to be directed in a feedforward direction (Bosman et al., 2012; van Kerkoerle et al., 2014; Bastos et al., 2015). In contrast to previous suggestions (Bastos et al., 2012), these interactions would involve the feedforward tansmission of sparse, informative stimulus responses rather than errors.

Gamma-band activity as a coding regime. As a second proposal, Vinck and Bosman (2016) suggested that gamma synchronization could also play an important role in coordinating the interactions between distributed parts of V1 receiving related, and thereby redundant, visual inputs. Evidence for this, though a direct test has yet to be conducted, has been reviewed in Vinck and Bosman (2016). Recently, and in line with the hypothesis, Lowet, Roberts, *Peter*, Gips and De Weerd (2017) have demonstrated that synchronization is a dynamical process between distant V1 regions *receiving similar inputs*. To the degree that predictable inputs generate a similar drive and thereby similar local oscillation frequency, TWCO predicts stronger synchronization.

In an excitatory-inhibitory network, this could lead to coordinated inhibition and redundancy removal. Furthermore, small differences in input are converted from a

rate to a phase code: slightly stronger inputs spike earlier within the oscillation cycle (Lowet et al., 2017). This resonates well with evidence that gamma-modulated cells show stronger stimulus selectivity and that strongly driven cells fire early in the gamma cycle (Womelsdorf et al., 2012; Onorato et al., 2019; Vinck et al., 2010a). Notably, this provides additional evidence that gamma-synchronized responses may constitute a specific type of coding regime. Similarity and predictability are related concepts. Highly similar (but not identical) inputs, would yield similar output rates, that are not easy to discriminate in a downstream area. In contrast, gamma synchronization may convert such activity into efficient, sparse and, by virtue of phase differences, discriminable responses. In line with this idea, Markowitz et al. (2008) have demonstrated in vitro that only cells with similar firing rates exhibited gamma-band synchrony. They suggested that this allows a many-are-equal computation and indicates the coexistence of rate- and spike synchrony coding. Interestingly, these considerations result in the prediction that gamma-band synchrony may be relevant for fine but not coarse stimulus discrimination. Preliminary evidence supporting this prediction are the increased neuronal disciminability of specifically similar stimuli observed with gamma-band increases in macaque V1 (Hansen and Dragoi, 2011) and the observation that bees can perform coarse but not fine odor discriminations when odor-induced oscillations are blocked (Stopfer et al., 1997).

Gamma-band activity and plasticity. The repetition effects studied in Chapters 2 and 3 are short-term plasticity effects that span the order of seconds to minutes. These could be brought about by changes in synaptic strength (Wang et al., 2011; Bazhenov et al., 2005), which can strengthen gamma-band activity. In a high-SNR regime with very strong input drive, synchronization might be prevented if the cells are in such a depolarized state that oscillatory membrane potential fluctuations no longer influence spike timing. Synaptic depression could therefore increase synchrony, especially for strong inputs, and sharpen stimulus responses. Synaptic depression could be instrumental in converting inputs from a rate-based to a phase-of-firing based coding regime as predicted by TWCO and reported e.g. by Womelsdorf et al. (2012), Lowet et al. (2017), Onorato et al. (2019) and Vinck et al. (2010a). This may aid stimulus discriminability (Markowitz et al., 2008). Crucially, oscillation phase-based codes may be easily learned by downstream networks, using spike timing-dependent plasticity mechanisms (Masquelier et al., 2009). In other words, rapid plasticity mechanisms may aid in the generation of gamma-band synchrony, which may be conducive to more long-term forms of learning.

Criticism. Criticisms of claims of a functional role for gamma-band oscillations quickly followed initial proposals (Shadlen and Movshon, 1999). Here, the questions of the causal role and volatility of gamma-band synchronization are briefly discussed. More extensive discussions that also involve considerations of effect size or the general problem of neuronal coordination can be found in, for example, Fries (2015); Singer (2018); Jia et al. (2013a,b); Burns et al. (2011); Merker (2013); Rolls et al. (2012). A major ongoing point of debate is the question of a causal role versus epiphenomenal quality of gamma-band activity. Does gamma-band synchronization implement computational functions, for example by a phase-of-firing code, or is it merely a signature of functional processes, such as interactions between excitatory and inhibitory cells, without any functional role?

Questions of causality are difficult to address in complex systems. Since gamma-band activity is a network phenomenon, there are many ways to influence gamma synchronization. However, such manipulations will inevitably influence other network properties, such as neuromodulation or the function of a particular cell type.

In an ideal scenario, an identical, subtle change in gamma-band oscillation behavior would be brought about through very different mechanisms, but lead to an identical functional consequence. Ideally, this would involve both the gain and loss of function through bidirectional manipulation of the different mechanisms. This would require an excellent understanding and fine-grained manipulation of the underlying circuit mechanisms. Until this challenge is mastered (if ever), it should be noted that gamma-band activity has influences on the membrane potential (Atallah and Scanziani, 2009), and thereby on spiking behavior, whose causal role is rarely questioned. Since oscillatory activity in cortex varies between areas and states and is therefore not unavoidable (Buzsáki and Wang, 2012; Buzsáki, 2006), and V1 gamma-band activity seems to involve specialized cell types that are highly informative about visual stimuli (Onorato et al., 2019; Gray and McCormick, 1996), it is at least possible that gamma-band activity plays a functional role by coordinating neuronal activity.

Another question regards the plausibility of a functional role for gamma-band activity given its volatile behavior. Specifically, gamma-band oscillations appear stochastic, with strong fluctuations in amplitude and frequency (Atallah and Scanziani, 2009; Spyropoulos et al., 2019; Lowet et al., 2017; Burns et al., 2011; Ray and Maunsell, 2010). However, recent experimental and modeling work has shown that the apparent stochasticity of gamma-band activity can be understood from the perspective of an ongoing synchronization process (Lowet et al., 2017). Furthermore, modeling work suggests that interareal neuronal networks can 1) spontaneously generate brief gamma-bursts which are 2) matched between areas, and 3) selectively route information flow in a flexible manner, since small modulation can switch between different routing states (Palmigiano et al., 2017). Frequency- and amplitude variation can enable synchronization. In summary, some aspects of gamma-band activity that could be considered problematic for function have the potential, upon closer examination, to reveal functional processing.

6.6 Methodological considerations and the psychology of technology

"It will have occured to the reader [...] that the science [...] must be in a very primitive condition, if it is still concerned with clarifying its fundamental notions at the stage we have been discussing. [...] This may be partly due to the sheer magnitude of the programme, [...], partly due to an intuitive confidence, [...], that, by amassing sufficient statistical material, all difficulties may ultimately be overcome."

R. A. Fisher (1936), "The Coefficient of Racial Likeness" and the Future of Craniometry.

Visual systems neuroscience, based on electrophysiology and particularly of the early visual system, provides a particular perspective on understanding vision that comes with specific limitations and biases. Olshausen and Field (2005) list some of the major limitations in the field, including biased sampling, biased stimuli, biased theory, and ignoring contextual influences and dynamics. How does this apply to the experiments in the preceding chapters?

Biased sampling. Visual neuroscience has been relying on recordings from single electrodes, inserted anew for each recording, for much of its existence. This has biased the recordings toward neurons that are visually responsive (also biasing stimuli to make responsiveness more likely). The use of chronically implanted recording arrays

with many electrodes, as in all experiments here, helps against this particular bias. The analysis of data that is not based on highly responsive cells brings new perspective, but also its own challenges (as discussed in Chapter 2). Moreover, chronic recording methods often do not provide single-cell resolution long-term, though this problem might be solved in various ways, in particular with miniaturization (Hong and Lieber, 2019). Even if this is solved, electrophysiology is still biased towards larger cells (since they are easier to detect) and particular layers (e.g. layer 1 is sparsely populated and right at the surface, i.e. easily damaged, Buzsáki et al. 2012). Newly developed imaging methods that can record from thousands of neurons regardless of their activity provide new perspectives on cortical function (e.g. Stosiek et al. 2003). Unfortunately, these come with their own biases, such as poor temporal resolution and bias towards upper layers, and the methods that make imaging possible might alter properties of neuronal responses themselves (McMahon and Jackson, 2018; Stosiek et al., 2003).

Biased stimuli. Technological constraints, such as attempts to detect responsive single neurons, can lead to biased stimuli. As an example, in the field of stimulus repetition, the advent of multielectrode recording arrays that covered wider parts of the visual field encouraged the use of larger stimuli. This change revealed both previously unknown coding principles, and similarities between areas that had been thought to engage in different computations with respect to repeated or prolonged stimulation (Wissig and Kohn, 2012; Solomon and Kohn, 2014). On the other hand, one could say that a desire to make optimal use of such arrays induces a bias towards larger stimuli, and in attempt to treat different sites the same, homogeneous ones. This choice has dramatic effects on contextual processes such as surround modulation and stimulus repetition effects, both of which are linked to gamma-band oscillations (Chapters 2 and 4).

Such observations have led to increasing attempts to study "natural" stimulation, typically in the form of achromatic still images (photographs). This paradigm shift, together with increasingly parallel recordings, has led to important new insights, and emphasized the power of theories inspired by optimal encoding of natural statistics or particular goals, such as object recognition, compared to ever-more detailed descriptive models (Schwartz and Simoncelli, 2001; Schwartz et al., 2017; Coen-Cagli et al., 2015; Stringer et al., 2019; DiCarlo et al., 2012). However, our world is neither static nor black-and-white. As seen in Chapters 2, 4 and 5, introducing color can have strong effects on both rate responses and gamma-band oscillations. Together with the ethological value of color information (Osorio and Vorobyev, 1996; Santos et al., 2001; Waitt et al., 2006; Gerald et al., 2007; Bichot et al., 2005; Corso et al., 2016; Melin et al., 2017), this suggests that a full understanding of primate vision will have to integrate color processing in efficient or predictive coding models (Wachtler et al., 2003). Similar arguments can be made for using naturalistic moving stimuli and self-motion (Hénaff et al., 2019; Britten, 2008).

The preceding chapters illustrated that using more naturalistic stimuli brings new insights but also new challenges (Chapter 2, Uran, Peter et al., in preparation). Generating controlled stimuli does not come without its own reasoning and advantages (Rust and Movshon, 2005). For example, a good theory is falsifiable with even a single carefully constructed counter-example. Due to the dependencies of different stimulus aspects in natural images (Barlow, 1959; Marr, 1982), this appears easier with artificial stimuli. Similarly, an artificial stimulus that provides a strong manipulation in line with a hypothesis can provide at least a quick initial test that can complement and be further elaborated with more extensive experiments using natural images (e.g. Chapter 4 provides an initial test for the "predictability" hypothesis for naturalistic input,

Uran, Peter et al., in preparation). Artificial stimulation can isolate aspects that tend to co-vary in natural stimulation, such as the red-green and blue-yellow color axes (Chapter 5). It can also be a powerful confirmation of models derived from natural stimulation (Bashivan et al., 2019). Clearly, an interactive use of both more naturalistic and artificial stimulation will be most informative (Stringer et al., 2019; Olshausen and Field, 2005; Rust and Movshon, 2005; Bashivan et al., 2019).

Biased theories. Theories tend to be strongly inspired by observable phenomena, which result from stimulus choice and the technology in use. In case of isolated recordings from a single cell, this restricts theories largely to rate-based encoding. Simultaneous recordings brought about an increasing recognition for the role of neuronal correlations, coordination and dynamics (Kohn et al., 2016; Solomon and Kohn, 2014; Gray et al., 1989; Singer, 2018; Fries, 2015; Olshausen and Field, 2005; Muller et al., 2018). Similarly, whole-brain imaging methods like fMRI were powerful in emphasizing the large network that underlies any cortical function or behavior. At the same time, these technologies de-emphasized the less visible aspects of cortical function, such as fast dynamics and input- versus output selectivity of a system (Sawamura et al., 2006; Logothetis and Wandell, 2004). It is always tempting to look for answers based on questions that can currently be tested, and challenging to be aware of blind spots of the current approach. Both the stimulus choice and change in perspective through new methods have been, and continue to be, transformative.

Approaches that can predict a relatively wide range of neuronal responses to natural stimuli currently are either based on artificial neuronal networks that are driven by large amounts of naturalistic training data and particular task goals (Kar et al., 2019; DiCarlo et al., 2012; Turner et al., 2019), or based on theories about the optimal encoding of naturalistic input (Schwartz and Simoncelli, 2001; Schwartz et al., 2017; Coen-Cagli et al., 2012; Barlow, 1959; Rao and Ballard, 1999). Since natural vision may involve goals both on the level of naturalistic tasks such as object recognition, and energy-efficient yet robust processing, the increasing integration of these approaches has the potential to move the field away from bias (Turner et al., 2019).

Contextual influences and dynamics. Arguably, all preceding Chapters have shown some awareness of this issue. For example, Chapters 2 and 3 have considered developments with a few stimulus repetitions, Chapter 4 tested spatial contextual influences on gamma-band responses, and the adaptation experiments in Chapter 5 were inspired by the observation of within-trial dynamics. Other important dynamical aspects have been beyond the scope of this work, including for example the dynamical interactions between different recording sites within a trial, the potential for traveling waves, the structuring of activity by rhythmic microsaccades, or a better understanding of the dynamics in the transient stimulus onset response (Lowet et al., 2017; Spyropoulos et al., 2019; Muller et al., 2018; Bosman et al., 2009; Lowet et al., 2016).

From a broader perspective, contextual influences are the core of efficient and predictive coding theories of visual processing (Rao and Ballard, 1999; Schwartz and Simoncelli, 2001; Schwartz et al., 2017; Coen-Cagli et al., 2012; Barlow, 1959). Recurrency, and thereby dynamics, are increasingly improving artificial neural network approaches (Kar et al., 2019; Turner et al., 2019; Kubilius et al., 2018; Quax and Van Gerven, 2018).

6.7 Conclusion and outlook

"The future is already here - it is just not very evenly distributed."
William Gibson.

The preceding chapters and discussion link gamma-band activity in V1 to contextual integration and recurrent processing as well as short-term plasticity. This thinking relates gamma-band activity to processes that can be described on different and equally valid levels. On one level of abstraction, gamma-band activity may reflect "predictability" in a high-SNR regime. This gives an intuitive unerstanding of the kind of stimuli that can generate strong gamma responses, and that might benefit the most from stimulus repetition. Moreover, "predictability" is a property that can be quantified for natural images in various ways, using for example algorithms that predict missing image information (Uran, Peter et al., in preparation). This could allow the generation of quantitative predictions for V1 responses to arbitrary stimuli.

"Predictability" by V1 results in surround suppression, which could be seen as the implementation of prediction in the circuit. Surround suppression requires sufficient drive and horizontal connectivity. From the perspective of gamma-band synchronization, this strongly links surround suppression to network dynamics, such as described using TWCO, where synchronization is the result of sufficient coupling (connectivity) and sufficiently similar inputs (drive). The dynamics reflect interactions in the V1 network, composed of various specialized cell types that may predispose the circuit for gamma-rhythmic stimulus processing. Network dynamics can be optimized by short-term plasticity through synaptic depression or facilitation, as is likely the case for stimulus repetition. Once established, a gamma-rhythmic phase-of-firing coding regime could result in consequences back on the functional level, such as increased fine discrimination.

An integration of efficient-coding and goal-oriented approaches to vision, incorporating contextual influences and recurrency, appears on the horizon. Such an approach links different levels of description, and given the strong dependence of V1 gamma on stimulus properties, context and recurrency, bears promise also for understanding gamma-band synchronization. The degree to which this promise can be fulfilled remains to be seen, and will most likely depend greatly on an understanding of the brain's powerful and plentiful mechanisms to learn and adapt.

Bibliography

Abbott, L. F. (1997). Synaptic depression and cortical gain control. *Science*, 275(5297):221–224.

Abeles, M. (1982). Role of the cortical neuron: integrator or coincidence detector? *Israel Journal of Medical Sciences*, 18:83.

Adjamian, P., Hadjipapas, A., Barnes, G. R., Hillebrand, A., and Holliday, I. E. (2008). Induced Gamma activity in primary visual cortex is related to luminance and not color contrast: An MEG study. *Journal of Vision*, 8(7):4.

Ahmed, B., Anderson, J. C., Douglas, R. J., Martin, K. A., and Whitteridge, D. (1998). Estimates of the net excitatory currents evoked by visual stimulation of identified neurons in cat visual cortex. *Cerebral Cortex*, 8(5):462–476.

Akam, T. and Kullmann, D. M. (2010). Oscillations and filtering networks support flexible routing of information. *Neuron*, 67(2):308–320.

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

Albrecht, D. G., Farrar, S. B., and Hamilton, D. B. (1984). Spatial contrast adaptation characteristics of neurones recorded in the cat's visual cortex. *Journal of Physiology*, 347(1):713–739.

Albright, T. D. and Stoner, G. R. (2002). Contextual influences on visual processing. *Annual Review of Neuroscience*, 25(1):339–379.

Alitto, H. J. and Usrey, W. M. (2008). Origin and dynamics of extraclassical suppression in the lateral geniculate nucleus of the macaque monkey. *Neuron*, 57(1):135–146.

Amado, C. and Kovács, G. (2016). Does surprise enhancement or repetition suppression explain visual mismatch negativity? *European Journal of Neuroscience*, 43(12):1590–1600.

Andersen, N., Krauth, N., and Nabavi, S. (2017). Hebbian plasticity in vivo: relevance and induction. *Current Opinion in Neurobiology*, 45:188–192.

Andolina, I. M., Jones, H. E., and Sillito, A. M. (2012). Effects of cortical feedback on the spatial properties of relay cells in the lateral geniculate nucleus. *Journal of Neurophysiology*, 109(3):889–899.

Angelucci, A., Bijanzadeh, M., Nurminen, L., Federer, F., Merlin, S., and Bressloff, P. C. (2017). Circuits and mechanisms for surround modulation in visual cortex. *Annual Review of Neuroscience*, 40:425–451.

Arnal, L. H. and Giraud, A. L. (2012). Cortical oscillations and sensory predictions. *Trends in Cognitive Sciences*, 16(7):390–398.

Arnal, L. H., Wyart, V., and Giraud, A.-L. (2011). Transitions in neural oscillations reflect prediction errors generated in audiovisual speech. *Nature Neuroscience*, 14(6):797–801.

Árni, K. and Campana, G. (2010). Where perception meets memory: A review of repetition priming in visual search tasks. *Attention, Perception, and Psychophysics*, 72(1):5–18.

Aschner, A., Solomon, S. G., Landy, M. S., Heeger, D. J., and Kohn, A. (2018). Temporal contingencies determine whether adaptation strengthens or weakens normalization. *Journal of Neuroscience*, 38(47):10129–10142.

Atallah, B. V. and Scanziani, M. (2009). Instantaneous modulation of gamma oscillation frequency by balancing excitation with inhibition. *Neuron*, 62:566–77.

Auksztulewicz, R. and Friston, K. (2016). Repetition suppression and its contextual determinants in predictive coding. *Cortex*, 80:125–140.

Axmacher, N., Mormann, F., Fernández, G., Elger, C. E., and Fell, J. (2006). Memory formation by neuronal synchronization. *Brain Research Reviews*, 52(1):170–182.

Azouz, R. and Gray, C. M. (2000). Dynamic spike threshold reveals a mechanism for synaptic coincidence detection in cortical neurons in vivo. *Proceedings of the National Academy of Sciences*, 97:8110–5.

Azouz, R. and Gray, C. M. (2003). Adaptive coincidence detection and dynamic gain control in visual cortical neurons in vivo. *Neuron*, 37:513–23.

Baker, C. I., Behrmann, M., and Olson, C. R. (2002). Impact of learning on representation of parts and wholes in monkey inferotemporal cortex. *Nature Neuroscience*, 5(11):1210–1216.

Baker, M. (2013). Neuroscience: Through the eyes of a mouse. *Nature*, 502(7470):156–158.

Bao, M. and Engel, S. A. (2012). Distinct mechanism for long-term contrast adaptation. *Proceedings of the National Academy of Sciences*, 109(15):5898–5903.

Barlow, H. (1959). Sensory mechanisms, the reduction of redundancy, and intelligence. *NPL Symposium on the Mechanization of Thought Process*, 10:535–539.

Barlow, H. (2001). Redundancy reduction revisited. *Network: Computation in Neural Systems*, 12(3):241–253.

Barone, P., Batardière, A., Knoblauch, K., and Kennedy, H. (2000). Laminar distribution of neurons in extrastriate areas projecting to visual areas V1 and V4 correlates with the hierarchical rank and indicates the operation of a distance rule. *Journal of Neuroscience*, 20(9):3263–3281.

Bartoli, E., Bosking, W., Li, Y., Beauchamp, M. S., Yoshor, D., and Foster, B. (2019). Distinct narrow and broadband gamma responses in human visual cortex. *bioRxiv*, page 572313.

Bartolo, M. J., Gieselmann, M. A., Vuksanovic, V., Hunter, D., Sun, L., Chen, X., Delicato, L. S., and Thiele, A. (2011). Stimulus-induced dissociation of neuronal firing rates and local field potential gamma power and its relationship to the resonance blood oxygen level-dependent signal in macaque primary visual cortex. *European Journal of Neuroscience*, 34(11):1857–1870.

Bartos, M., Vida, I., and Jonas, P. (2007). Synaptic mechanisms of synchronized gamma oscillations in inhibitory interneuron networks. *Nature Reviews Neuroscience*, 8:45–56.

Bashivan, P., Kar, K., and DiCarlo, J. J. (2019). Neural population control via deep image synthesis. *Science*, 364(6439).

Bastos, A. M., Briggs, F., Alitto, H. J., Mangun, G. R., and Usrey, W. M. (2014). Simultaneous recordings from the primary visual cortex and lateral geniculate nucleus reveal rhythmic interactions and a cortical source for gamma-band oscillations. *Journal of Neuroscience*, 34(22):7639–7644.

Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., and Friston, K. J. (2012). Canonical microcircuits for predictive coding. *Neuron*, 76(4):695–711.

Bastos, A. M., Vezoli, J., Bosman, C. A., Schoffelen, J.-M., Oostenveld, R., Dowdall, J. R., De Weerd, P., Kennedy, H., and Fries, P. (2015). Visual areas exert feedforward and feedback influences through distinct frequency channels. *Neuron*, 85(2):390–401.

Bazhenov, M., Stopfer, M., Sejnowski, T. J., and Laurent, G. (2005). Fast odor learning improves reliability of odor responses in the locust antennal lobe. *Neuron*, 46(3):483–492.

Berkes, P., Orbán, G., Lengyel, M., and Fiser, J. (2011). Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science*, 331(6013):83–87.

Bernander, Ö., Koch, C., and Usher, M. (1994). The effect of synchronized inputs at the single neuron level. *Neural Computation*, 6:622–641.

Bernstein, I. S. (2005). Macaque societies. A model for the study of social organization. *International Journal of Primatology*, 26(5):1199–1201.

Bessaih, T., Higley, M. J., and Contreras, D. (2018). Millisecond precision temporal encoding of stimulus features during cortically generated gamma oscillations in the rat somatosensory cortex. *Journal of Physiology*, 596(3):515–534.

Besserve, M., Lowe, S. C., Logothetis, N. K., Schölkopf, B., and Panzeri, S. (2015a). Shifts of gamma phase across primary visual cortical sites reflect dynamic stimulus-modulated information transfer. *PLoS Biology*, 13(9):e1002257.

Besserve, M., Lowe, S. C., Logothetis, N. K., Schölkopf, B., and Panzeri, S. (2015b). Shifts of gamma phase across primary visual cortical sites reflect dynamic stimulus-modulated information transfer. *PLoS Biology*, 13(9):e1002257.

Bichot, N. P., Rossi, A. F., and Desimone, R. (2005). Parallel and serial neural mechanisms for visual search in macaque area V4. *Science*, 308:529–34.

Bliss, T. V. and Lomo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *Journal of Physiology*, 232:331–356.

Börgers, C. and Kopell, N. (2003). Synchronization in networks of excitatory and inhibitory neurons with sparse, random connectivity. *Neural Computation*, 15(3):509–538.

Börgers, C. and Kopell, N. (2005). Effects of noisy drive on rhythms in networks of excitatory and inhibitory neurons. *Neural Computation*, 17:557–608.

Börgers, C. and Kopell, N. J. (2008). Gamma oscillations and stimulus selection. *Neural Computation*, 20:383–414.

Bosman, C., Schoffelen, J., Brunet, N., Oostenveld, R., Bastos, A., Womelsdorf, T., Rubehn, B., Stieglitz, T., De Weerd, P., and Fries, P. (2012). Attentional stimulus selection through selective synchronization between monkey visual areas. *Neuron*, 75:875–888.

Bosman, C. A., Womelsdorf, T., Desimone, R., and Fries, P. (2009). A microsaccadic rhythm modulates gamma-band synchronization and behavior. *Journal of Neuroscience*, 29:9471–80.

Boudreau, C. E. (2005). Short-term depression in thalamocortical synapses of cat primary visual cortex. *Journal of Neuroscience*, 25(31):7179–7190.

Bradley, M. M., Miccoli, L., Escrig, M. A., and Lang, P. J. (2008). The pupil as a measure of emotional arousal and autonomic activation. *Psychophysiology*, 45(4):602–607.

Brady, T. F., Konkle, T., Alvarez, G. A., and Oliva, A. (2008). Visual long-term memory has a massive storage capacity for object details. *Proceedings of the National Academy of Sciences*, 105(38):14325–14329.

Brainard, D. (1996). Cone contrast and opponent modulation color spaces. In *Human color vision*. Opt. Soc. Am.

Bressler, S. L., Coppola, R., and Nakamura, R. (1993). Episodic multiregional cortical coherence at multiple frequencies during visual task performance. *Nature*, 366:153–156.

Bressler, S. L., Richter, C. G., Chen, Y., and Ding, M. (2006). Top-down cortical influences in visual expectation. In *International Joint Conference on Neural Networks*, pages 188–194.

Britten, K. H. (2008). Mechanisms of self-motion perception. *Annual Review of Neuroscience*, 31(1):389–410.

Brodmann, K. (1909). *Vergleichende Lokalisationslehre der Großhirnrinde*. Barth, Leipzig.

Brunet, N., Bosman, C. A., Roberts, M., Oostenveld, R., Womelsdorf, T., De Weerd, P., and Fries, P. (2015). Visual cortical gamma-band activity during free viewing of natural images. *Cerebral Cortex*, 25(4):918–926.

Brunet, N. M., Bosman, C. A., Vinck, M., Roberts, M., Oostenveld, R., Desimone, R., De Weerd, P., and Fries, P. (2014). Stimulus repetition modulates gamma-band synchronization in primate visual cortex. *Proceedings of the National Academy of Sciences*, 111(9):3626–3631.

Brunet, N. M. and Fries, P. (2019). Human visual cortical gamma reflects natural image structure. *NeuroImage*, 200:635–643.

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

Burns, S. P., Xing, D., and Shapley, R. M. (2011). Is gamma-band activity in the local field potential of V1 cortex a "clock" or filtered noise? *Journal of Neuroscience*, 31(26):9658–9664.

Buschman, T. J. and Miller, E. K. (2007). Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. *Science*, 315:1860–2.

Bush, P. and Sejnowski, T. (1996). Inhibition synchronizes sparsely connected cortical neurons within and between columns in realistic network models. *Journal of Computational Neuroscience*, 3(2):91–110.

Buzsáki, G. (2006). *Rhythms of the Brain*. Oxford University Press, USA.

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

Buzsáki, G. and Draguhn, A. (2004). Neuronal oscillations in cortical networks. *Science*, 304:1926–1929.

Buzsáki, G., Logothetis, N., and Singer, W. (2013). Scaling brain size, keeping timing: Evolutionary preservation of brain rhythms. *Neuron*, 80(3):751–764.

Buzsáki, G. and Wang, X. J. (2012). Mechanisms of gamma oscillations. *Annual Review of Neuroscience*, 35:203–225.

Cajal, S. R. Y. (1894). The Croonian Lecture: La fine structure des centres nerveux. *Proceedings of the Royal Society of London*, 55(331-335):444–468.

Calabrese, E., Badea, A., Coe, C. L., Lubach, G. R., Shi, Y., Styner, M. A., and Johnson, G. A. (2015). A diffusion tensor MRI atlas of the postmortem rhesus macaque brain. *NeuroImage*, 117:408–416.

Camp, A. J., Tailby, C., and Solomon, S. G. (2009). Adaptable mechanisms that regulate the contrast response of neurons in the primate lateral geniculate nucleus. *Journal of Neuroscience*, 29(15):5009–5021.

Caporale, N. and Dan, Y. (2008). Spike timing-dependent plasticity: A Hebbian learning rule. *Annual Review of Neuroscience*, 31(1):25–46.

Carandini, M., Demb, J. B., Mante, V., Tolhurst, D. J., Dan, Y., Olshausen, B. A., Gallant, J. L., and Rust, N. C. (2005). Do we know what the early visual system does? *Journal of Neuroscience*, 25(46):10577–97.

Carandini, M. and Heeger, D. J. (2011). Normalization as a canonical neural computationation. *Nature Reviews Neuroscience*.

Carandini, M. and Heeger, D. J. (2012). Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, 13(1):51–62.

Cardin, J. A., Carlén, M., Meletis, K., Knoblich, U., Zhang, F., Deisseroth, K., Tsai, L.-H., and Moore, C. I. (2009). Driving fast-spiking cells induces gamma rhythm and controls sensory responses. *Nature*, 459:663–7.

Cardin, J. A., Palmer, L. A., and Contreras, D. (2005). Stimulus-dependent gamma (30-50 Hz) oscillations in simple and complex fast rhythmic bursting cells in primary visual cortex. *Journal of Neuroscience*, 25:5339–50.

Castelo-Branco, M., Neuenschwander, S., and Singer, W. (1998). Synchronization of visual responses between the cortex, lateral geniculate nucleus, and retina in the anesthetized cat. *Journal of Neuroscience*, 18(16):6395–6410.

Castro-Alamancos, M. A. (2002). Role of thalamocortical sensory suppression during arousal: Focusing sensory inputs in neocortex. *Journal of Neuroscience*, 22(22):9651–9655.

Cavanaugh, J. R., Bair, W., and Movshon, J. A. (2002). Selectivity and spatial distribution of signals from the receptive field surround in macaque V1 neurons. *Journal of Neurophysiology*, 88(5):2547–2556.

Chalk, M., Gutkin, B., and Denève, S. (2016). Neural oscillations as a signature of efficient coding in the presence of synaptic delays. *eLife*, 5.

Chalk, M., Herrero, J. L., Gieselmann, M. A., Delicato, L. S., Gotthardt, S., and Thiele, A. (2010). Attention reduces stimulus-driven gamma frequency oscillations and spike field coherence in V1. *Neuron*, 66(1):114–125.

Chang, L., Bao, P., and Tsao, D. Y. (2017). The representation of colored objects in macaque color patches. *Nature Communications*, 8(1).

Changizi, M. A., Zhang, Q., and Shimojo, S. (2006). Bare skin, blood and the evolution of primate colour vision. *Biology Letters*, 2(2):217–221.

Chao, Z. C., Takaura, K., Wang, L., Fujii, N., and Dehaene, S. (2018). Large-scale cortical networks for hierarchical prediction and prediction error in the primate brain. *Neuron*, 100(5):1252–1266.e3.

Chaparro, A., Huang, E., Kronauer, R., Eskew, R. T., et al. (1993). Colour is what the eye sees best. *Nature*, 361(6410):348.

Chen, C. M., Lakatos, P., Shah, A. S., Mehta, A. D., Givre, S. J., Javitt, D. C., and Schroeder, C. E. (2007). Functional anatomy and interaction of fast and slow visual pathways in macaque monkeys. *Cerebral Cortex*, 17(7):1561–1569.

Chen, G., Zhang, Y., Li, X., Zhao, X., Ye, Q., Lin, Y., Tao, H. W., Rasch, M. J., and Zhang, X. (2017). Distinct inhibitory circuits orchestrate cortical beta and gamma band oscillations. *Neuron*, 96(6):1403–1418.

Cheung, B., Weiss, E., and Olshausen, B. A. (2016). Emergence of foveal image sampling from learning to attend in visual scenes. *CoRR*, abs/1611.09430.

Coen-Cagli, R., Dayan, P., and Schwartz, O. (2012). Cortical surround interactions and perceptual salience via natural scene statistics. *PLoS Computational Biology*, 8(3):e1002405.

Coen-Cagli, R., Kohn, A., and Schwartz, O. (2015). Flexible gating of contextual influences in natural vision. *Nature Neuroscience*, 18(11):1648–55.

Colgin, L., Denninger, T., Fyhn, M., Hafting, T., Bonnevie, T., Jensen, O., Moser, M., and Moser, E. (2009). Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature*, 462:353–357.

Constantinidis, C., Funahashi, S., Lee, D., Murray, J. D., Qi, X. L., Wang, M., and Arnsten, A. F. (2018). Persistent spiking activity underlies working memory. *Journal of Neuroscience*, 38(32):7020–7028.

Conway, B. R. and Tsao, D. Y. (2006). Color architecture in alert macaque cortex revealed by fMRI. *Cerebral Cortex*, 16(11):1604–1613.

Cooke, S. and Bear, M. F. (2012). Stimulus-selective response plasticity in the visual cortex: An assay for the assessment of pathophysiology and treatment of cognitive impairment associated with psychiatric disorders. *Biological Psychiatry*, 71(6):487–495.

Cooke, S. F. and Bear, M. F. (2010). Visual experience induces long-term potentiation in the primary visual cortex. *Journal of Neuroscience*, 30(48):16304–16313.

Cooke, S. F. and Bear, M. F. (2014). How the mechanisms of long-term synaptic potentiation and depression serve experience-dependent plasticity in primary visual cortex. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1633).

Cooke, S. F., Komorowski, R. W., Kaplan, E. S., Gavornik, J. P., and Bear, M. F. (2015). Visual recognition memory, manifested as long-term habituation, requires synaptic plasticity in V1. *Nature Neuroscience*, 18(2):262–271.

Corso, J., Bowler, M., Heymann, E. W., Roos, C., and Mundy, N. I. (2016). Highly polymorphic colour vision in a new world monkey with red facial skin, the bald uakari (Cacajao calvus). *Proceedings of the Royal Society B: Biological Sciences*, 283(1828).

Cumming, B. G. and DeAngelis, G. C. (2001). The physiology of stereopsis. *Annual Review of Neuroscience*, 24:203–38.

Dasgupta, S., Stevens, C. F., and Navlakha, S. (2017). A neural algorithm for a fundamental computing problem. *Science*, 358(6364):793–796.

De Baene, W. and Vogels, R. (2010). Effects of adaptation on the stimulus selectivity of macaque inferior temporal spiking activity and local field potentials. *Cerebral Cortex*, 20(9):2145–2165.

De Lange, F. P., Heilbron, M., and Kok, P. (2018). How do expectations shape perception? *Trends in Cognitive Sciences*, 22(9):764–779.

De Valois, R. L. (1965). Analysis and coding of color vision in the primate visual system. *Cold Spring Harbor Symposia on Quantitative Biology*, 30:567–579.

De Valois, R. L., Cottaris, N. P., Elfar, S. D., Mahon, L. E., and Wilson, J. A. (2002). Some transformations of color information from lateral geniculate nucleus to striate cortex. *Proceedings of the National Academy of Sciences*, 97(9):4997–5002.

De Valois, R. L., De Valois, K. K., and Mahon, L. E. (2000). Contribution of S opponent cells to color appearance. *Proceedings of the National Academy of Sciences*, 97(1):512–517.

De Valois, R. L., Morgan, H., and Snodderly, D. M. (1974a). Psychophysical studies of monkey Vision-III. Spatial luminance contrast sensitivity tests of macaque and human observers. *Vision Research*, 14(1):75–81.

De Valois, R. L., Morgan, H. C., Polson, M. C., Mead, W. R., and Hull, E. M. (1974b). Psychophysical studies of monkey vision-I. Macaque luminosity and color vision tests. *Vision Research*, 14(1):53–67.

De Weerd, P., Pinaud, R., and Bertini, G. (2006). Plasticity in V1 induced by perceptual learning. In *Plasticity in the Visual System*, pages 245–283. Springer.

DeFelipe, J. (2015). The dendritic spine story: an intriguing process of discovery. *Frontiers in Neuroanatomy*, 9.

Derrington, A. M., Krauskopf, J., and Lennie, P. (1984). Chromatic mechanisms in lateral geniculate nucleus of macaque. *Journal of Physiology*, 357(1):241–265.

Desimone, R. (1996). Neural mechanisms for visual memory and their role in attention. *Proceedings of the National Academy of Sciences*, 93(24):13494–13499.

Dhruv, N. T., Tailby, C., Sokol, S. H., and Lennie, P. (2011). Multiple adaptable mechanisms early in the primate visual pathway. *Journal of Neuroscience*, 31(42):15016–15025.

DiCarlo, J. J., Zoccolan, D., and Rust, N. C. (2012). How does the brain solve visual object recognition? *Neuron*, 73(3):415–434.

Douglas, R. J. and Martin, K. A. (1991). A functional microcircuit for cat visual cortex. *Journal of Physiology*, 440(1):735–769.

Douglas, R. J., Martin, K. A., and Whitteridge, D. (2008). A canonical microcircuit for neocortex. *Neural Computation*, 1(4):480–488.

Dowdall, J., Schmiedt, J., Stephan, M., and Fries, P. (2018). Arcade: a modular multithreaded stimulus presentation software for the real-time control of stimuli, actions and reward during behavioral experiments. *SfN 2018. 254.18.*

Dowdall, J., Scholvinck, M., Zhang, Y., Peter, A., Vezoli, J., and Fries, P. (2019). Rapid learning of a fixation-based visual selective attention task and its effects on V1 firing rates. *SfN 2019. 664.13.*

Dowling, J. E. (1970). Organization of vertebrate retinas. *Investigative Ophthalmology*, 9(9):655–80.

Dragoi, V., Sharma, J., Miller, E. K., and Sur, M. (2002). Dynamics of neuronal sensitivity in visual cortex and local feature discrimination. *Nature Neuroscience*, 5(9):883–891.

Dubuc, C., Allen, W. L., Maestripieri, D., and Higham, J. P. (2014). Is male rhesus macaque red color ornamentation attractive to females? *Behavioral Ecology and Sociobiology*, 68(7):1215–1224.

D'Zmura, M. (1991). Color in visual search. *Vision Research*, 31(6):951–966.

Ebitz, R. B., Pearson, J. M., and Platt, M. L. (2014). Pupil size and social vigilance in rhesus macaques. *Frontiers in Neuroscience*, 8:100.

Edelman, J. A. and Keller, E. L. (1996). Activity of visuomotor burst neurons in the superior colliculus accompanying express saccades. *Journal of Neurophysiology*, 76(2):908–26.

Ekstrom, A. (2010). How and when the fMRI BOLD signal relates to underlying neural activity: The danger in dissociation. *Brain Research Reviews*, 62(2):233–244.

Engbert, R. and Kliegl, R. (2003). Microsaccades uncover the orientation of covert attention. *Vision Research*, 43(9):1035–1045.

Engel, A. K., Fries, P., and Singer, W. (2001). Dynamic predictions: oscillations and synchrony in top-down processing. *Nature Reviews Neuroscience*, 2:704–716.

Engel, A. K., König, P., Kreiter, A. K., Schillen, T. B., and Singer, W. (1992). Temporal coding in the visual cortex: new vistas on integration in the nervous system. *Trends in Neuroscience*, 15:218–26.

Eskandar, E. N., Richmond, B. J., and Optican, L. M. (1992). Role of inferior temporal neurons in visual memory I. Temporal encoding of information about visual images, recalled images, and behavioral context. *Journal of Neurophysiology*, 68(4):1277–1295.

Fabbrini, F., Van den Haute, C., De Vitis, M., Baekelandt, V., Vanduffel, W., and Vogels, R. (2019). Probing the mechanisms of repetition suppression in inferior temporal cortex with optogenetics. *Current Biology*, 29(12):1988–1998.

Felleman, D. (2004). Visual System in the Brain. In *International Encyclopedia of the Social & Behavioral Sciences*, pages 16278–16285.

Felleman, D. and Van Essen, D. (1991). Distributed hierarchical processing in primate visual cortex. *Cerebral cortex*, 1:1–47.

Field, D. J. (1987). Relations between the statistics of natural images and the response properties of cortical cells. *Journal of the Optical Society of America A*, 4(12):2379.

Fioravante, D. and Regehr, W. G. (2011). Short-term forms of presynaptic plasticity. *Current Opinion in Neurobiology*, 21(2):269 – 274.

Fiser, J., Berkes, P., Orbán, G., and Lengyel, M. (2010). Statistically optimal perception and learning: from behavior to neural representations. *Trends in Cognitive Sciences*, 14(3):119–30.

Fisher, R. A. (1936). "The Coefficient of Racial Likeness" and the future of craniometry. *Journal of the Royal Anthropological Institute of Great Britain and Ireland*, 66:57.

Foster, K. W., Saranak, J., Patel, N., Zarilli, G., Okabe, M., Kline, T., and Nakanishi, K. (1984). A rhodopsin is the functional photoreceptor for phototaxis in the unicellular eukaryote Chlamydomonas. *Nature*, 311(5988):756–759.

Fourier, J. (1822). Théorie analytique de la chaleur (1822). In Grattan-Guinness, I., Cooke, R., Corry, L., Crépel, P., and Guicciardini, N., editors, *Landmark Writings in Western Mathematics 1640-1940*, pages 354 – 365. Elsevier Science, Amsterdam.

Fournier, J., Monier, C., Pananceau, M., and Frégnac, Y. (2011). Adaptation of the simple or complex nature of V1 receptive fields to visual statistics. *Nature Neuroscience*, 14(8):1053–1060.

Freedman, D. J., Riesenhuber, M., Poggio, T., and Miller, E. K. (2006). Experience-dependent sharpening of visual shape selectivity in inferior temporal cortex. *Cerebral Cortex*, 16(11):1631–1644.

Freeman, J., Ziemba, C. M., Heeger, D. J., Simoncelli, E. P., and Movshon, J. A. (2013). A functional and perceptual signature of the second visual area in primates. *Nature Neuroscience*, 16(7):974–981.

Freeman, W. J. (2007). Definitions of state variables and state space for brain-computer interface. *Cognitive Neurodynamics*, 1(2):85–96.

Frenkel, M. Y., Sawtell, N. B., Diogo, A. C. M., Yoon, B., Neve, R. L., and Bear, M. F. (2006). Instructive effect of visual experience in mouse visual cortex. *Neuron*, 51(3):339–349.

Friedman, H. S., Zhou, H., and von der Heydt, R. (2003). The coding of uniform colour figures in monkey visual cortex. *Journal of Physiology*, 548(2):593–613.

Frien, A., Eckhorn, R., Bauer, R., Woelbern, T., and Gabriel, A. (2000). Fast oscillations display sharper orientation tuning than slower components of the same recordings in striate cortex of the awake monkey. *European Journal of Neuroscience*, 12(4):1453–1465.

Fries, P. (2005). A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends in Cognitive Sciences*, 9:474–480.

Fries, P. (2009). Neuronal gamma-band synchronization as a fundamental process in cortical computation. *Annual Review of Neuroscience*, 32:209–224.

Fries, P. (2015). Rhythms for cognition: communication through coherence. *Neuron*, 88(1):220–235.

Fries, P., Bastos, A. M., Mangun, G. R., Adams, R. A., Friston, K. J., and Usrey, W. M. (2012). Canonical microcircuits for predictive coding. *Neuron*, 76(4):695–711.

Fries, P., Nikolic, D., and Singer, W. (2007). The gamma cycle. *Trends in Neurosciences*, 30:309–316.

Friese, U., Supp, G. G., Hipp, J. F., Engel, A. K., and Gruber, T. (2012). Oscillatory MEG gamma band activity dissociates perceptual and conceptual aspects of visual object processing: A combined repetition/conceptual priming study. *NeuroImage*, 59(1):861–871.

Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1456):815–836.

Friston, K. (2008). Hierarchical models in the brain. *PLoS Computational Biology*, 4(11):e1000211.

Friston, K. (2012). The history of the future of the Bayesian brain. *NeuroImage*, 62(2):1230–1233.

Fuster, J. M. and Alexander, G. E. (1971). Neuron activity related to short-term memory. *Science*, 173(3997):652–654.

Gagin, G., Bohon, K. S., Butensky, A., Gates, M. A., Hu, J.-Y., Lafer-Sousa, R., Pulumo, R. L., Qu, J., Stoughton, C. M., Swanbeck, S. N., and Conway, B. R. (2014). Color-detection thresholds in rhesus macaque monkeys and humans. *Journal of Vision*, 14(8):12–12.

Gail, A. (2000). Contour decouples gamma activity across texture representation in monkey striate cortex. *Cerebral Cortex*, 10(9):840–850.

Galuske, R. A. W., Munk, M. H. J., and Singer, W. (2019). Relation between gamma oscillations and neuronal plasticity in the visual cortex. accepted for publication. *PNAS*.

Ganmor, E., Katz, Y., and Lampl, I. (2010). Intensity-dependent adaptation of cortical and thalamic neurons is controlled by brainstem circuits of the sensory pathway. *Neuron*, 66(2):273–286.

Garrido, M. I., Kilner, J. M., Kiebel, S. J., Stephan, K. E., Baldeweg, T., and Friston, K. J. (2009). Repetition suppression and plasticity in the human brain. *NeuroImage*, 48(1):269–279.

Gattass, R., Gross, C. G., and Sandell, J. H. (1981). Visual topography of V2 in the macaque. *Journal of Comparative Neurology*, 201(4):519–539.

Gattass, R., Lima, B., Soares, J. G. M., and Ungerleider, L. G. (2015). Controversies about the visual areas located at the anterior border of area V2 in primates. *Visual Neuroscience*, 32.

Gattass, R., Sousa, A. P., and Rosa, M. G. (1987). Visual topography of V1 in the Cebus monkey. *Journal of Comparative Neurology*, 259(4):529–548.

Gerald, M. S., Waitt, C., Little, A. C., and Kraiselburd, E. (2007). Females pay attention to female secondary sexual color: An experimental study in Macaca mulatta. *International Journal of Primatology*, 28(1):1–7.

Ghodrati, M., Morris, A. P., and Price, N. S. C. (2015). The (un)suitability of modern liquid crystal displays (LCDs) for vision research. *Frontiers in Psychology*, 6(303):1–11.

Ghuman, A. S., Bar, M., Dobbins, I. G., and Schnyer, D. M. (2008). The effects of priming on frontal-temporal communication. *Proceedings of the National Academy of Sciences*, 105(24):8405–8409.

Giaschi, D., Douglas, R., Marlin, S., and Cynader, M. (1993). The time course of direction-selective adaptation in simple and complex cells in cat striate cortex. *Journal of Neurophysiology*, 70(5):2024–2034.

Gieselmann, M. A. and Thiele, A. (2008). Comparison of spatial integration and surround suppression characteristics in spiking activity and the local field potential in macaque V1. *European Journal of Neuroscience*, 28:447–459.

Gil, Z., Connors, B. W., and Amitai, Y. (1997). Differential regulation of neocortical synapses by neuromodulators and activity. *Neuron*, 19(3):679–686.

Gilbert, C. (2013). Contextual modulation in the visual cortex. *Journal of Vision*, 13(9):1369–1369.

Gilbert, C. D. (1992). Horizontal integration and cortical dynamics. *Neuron*, 9(1):1–13.

Gilbert, J., Gotts, S., Carver, F., and Martin, A. (2010). Object repetition leads to local increases in the temporal coordination of neural responses. *Frontiers in Human Neuroscience*, 4:30.

Goldman-Rakic, P. S. (1995). Cellular basis of working memory. *Neuron*, 14(3):477–485.

Golgi, C. (1875). Sulla fina struttura del bulbi olfattorii. *Reggio-Emilia: Printer Stefano Calderini*.

Goodale, M. A. and Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends in Neurosciences*, 15(1):20–25.

Gotts, S. (2003). Mechanisms underlying enhanced processing efficiency in neural systems. *Dissertation*, pages 1–324.

Gotts, S. J., Chow, C. C., and Martin, A. (2012). Repetition priming and repetition suppression: A case for enhanced efficiency through neural synchronization. *Cognitive Neuroscience*, 3(3-4):227–237.

Gray, C. and McCormick, D. (1996). Chattering cells: superficial pyramidal neurons contributing to the generation of synchronous oscillations in the visual cortex. *Science*, 274:109.

Gray, C. M., Engel, A. K., König, P., and Singer, W. (1990). Stimulus-dependent neuronal oscillations in cat visual cortex: Receptive field properties and feature dependence. *European Journal of Neuroscience*, 2(7):607–619.

Gray, C. M., König, P., Engel, A. K., and Singer, W. (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature*, 338:334–337.

Gregoriou, G. G., Gotts, S. J., Zhou, H., and Desimone, R. (2009). High-frequency, long-range coupling between prefrontal and visual cortex during attention. *Science*, 324:1207–1210.

Grill-Spector, K., Henson, R., and Martin, A. (2006). Repetition and the brain: Neural models of stimulus-specific effects. *Trends in Cognitive Sciences*, 10(1):14–23.

Grosof, D. H., Shapley, R. M., and Hawken, M. J. (1993). Macaque V1 neurons can signal "illusory" contours. *Nature*, 365(6446):550–552.

Gross, C. G. (2002). Genealogy of the "grandmother cell". *Neuroscientist*, 8(5):512–518.

Grossberg, S. (1997). Are there universal principles of brain computation? In *Mathematics of Neural Networks*, pages 34–40. Springer.

Grossberg, S. (2013). Adaptive resonance theory: How a brain learns to consciously attend, learn, and recognize a changing world. *Neural Networks*, 37:1–47.

Grossberg, S. and Pearson, L. R. (2008). Laminar cortical dynamics of cognitive and motor working memory, sequence learning and performance: Toward a unified theory of how the cerebral cortex works. *Psychological Review*, 115(3):677–732.

Grothe, I., Neitzel, S. D., Mandon, S., and Kreiter, A. K. (2012). Switching neuronal inputs by differential modulations of gamma-band phase-coherence. *Journal of Neuroscience*, 32:16172–16180.

Güçlü, U. and van Gerven, M. A. J. (2015). Deep neural networks reveal a gradient in the complexity of neural representations across the ventral stream. *Journal of Neuroscience*, 35(27):10005–14.

Gutnisky, D. A. and Dragoi, V. (2008). Adaptive coding of visual information in neural populations. *Nature*, 452(7184):220–224.

Hadjipapas, A., Lowet, E., Roberts, M. J., Peter, A., and De Weerd, P. (2015). Parametric variation of gamma frequency and power with luminance contrast: A comparative study of human MEG and monkey LFP and spike responses. *NeuroImage*, 112:327–340.

Haider, B., Krause, M. R., Duque, A., Yu, Y., Touryan, J., Mazer, J. A., and McCormick, D. A. (2010). Synaptic and network mechanisms of sparse and reliable visual cortical activity during nonclassical receptive field stimulation. *Neuron*, 65(1):107–121.

Haller, M., Donoghue, T., Peterson, E., Varma, P., Sebastian, P., Gao, R., Noto, T., Knight, R. T., Shestyuk, A., and Voytek, B. (2018). Parameterizing neural power spectra. *bioRxiv*.

Hansen, B. J. and Dragoi, V. (2011). Adaptation-induced synchronization in laminar cortical circuits. *Proceedings of the National Academy of Sciences*, 108(26):10720–10725.

Hansen, T. and Gegenfurtner, K. R. (2013). Higher order color mechanisms: evidence from noise-masking experiments in cone contrast space. *Journal of Vision*, 13(1):26–26.

Hárosi, F. I. (1987). Cynomolgus and rhesus monkey visual pigments. application of fourier transform smoothing and statistical techniques to the determination of spectral parameters. *Journal of General Physiology*, 89(5):717–743.

Harris, K. D., Csicsvari, J., Hirase, H., Dragoi, G., and Buzsáki, G. (2003). Organization of cell assemblies in the hippocampus. *Nature*, 424(6948):552–6.

Harvey, B. M. and Dumoulin, S. O. (2011). Therelationship between cortical magnification factor and population receptive field size in human visual cortex: Constancies in cortical architecture. *Journal of Neuroscience*, 31(38):13604–13612.

Hasenstaub, A., Shu, Y., Haider, B., Kraushaar, U., Duque, A., and McCormick, D. A. (2005). Inhibitory postsynaptic potentials carry synchronized frequency information in active cortical networks. *Neuron*, 47:423–435.

Hattori, D., Aso, Y., Swartz, K. J., Rubin, G. M., Abbott, L. F., and Axel, R. (2017). Representations of novelty and familiarity in a mushroom body compartment. *Cell*, 169(5):956–969.

Havenith, M. N., Yu, S., Biederlack, J., Chen, N. H., Singer, W., and Nikolic, D. (2011). Synchrony makes neurons fire in sequence, and stimulus properties determine who is ahead. *Journal of Neuroscience*, 31:8570–8584.

Hawkins, J., George, D., and Niemasik, J. (2009). Sequence memory for prediction, inference and behaviour. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1521):1203–1209.

Hebb, D. O. (1949). *The organization of behavior: A neuropsychological theory*. New York: Wiley & Sons.

Heeger, D. J. (2017). Theory of cortical function. *Proceedings of the National Academy of Sciences*, 114(8):1773–1782.

Hénaff, O. J., Goris, R. L., and Simoncelli, E. P. (2019). Perceptual straightening of natural videos. *Nature Neuroscience*.

Henrie, J. and Shapley, R. (2005). LFP power spectra in V1 cortex: the graded effect of stimulus contrast. *Journal of Neurophysiology*, 94:479.

Henson, R. N. (2003). Neuroimaging studies of priming. *Progress in Neurobiology*, 70(1):53–81.

Henson, R. N. (2012). Repetition accelerates neural dynamics: In defense of facilitation models. *Cognitive Neuroscience*, 3(3-4):240–241.

Herculano-Houzel, S., Munk, M. H., Neuenschwander, S., and Singer, W. (1999). Precisely synchronized oscillatory firing patterns require electroencephalographic activation. *Journal of Neuroscience*, 19:3992–4010.

Hermes, D., Kay, K., and Winawer, J. (2014). Stimulus selectivity of broadband field potentials, but not gamma oscillations, matches population responses as measured by BOLD fMRI in human visual cortex. *Journal of Vision*, 14(10):1432–1432.

Hermes, D., Miller, K., Wandell, B., and Winawer, J. (2015). Stimulus dependence of gamma oscillations in human visual cortex. *Cerebral Cortex*, 25(9):2951–2959.

Hermes, D., Petridou, N., Kay, K., and Winawer, J. (2019). An image-computable model for the stimulus selectivity of gamma oscillations. *bioRxiv*.

Hesse, J. and Gross, T. (2014). Self-organized criticality as a fundamental property of neural systems. *Frontiers in Systems Neuroscience*, 8.

Hinton, G. E. (2000). Computation by neural networks. *Nature Neuroscience*, 3(11s):1170.

Ho, C. L. A., Zimmerman, R., Weidinger, J., Tia, B., Pfifferi, F., Aujard, F., Wolf, F., and Huber, D. (2019). Orientation preference maps in the mouse lemur confirm common design principles in primate visual cortex. *Cosyne Abstracts 2019, Lisbon, PT*.

Hong, G. and Lieber, C. M. (2019). Novel electrode technologies for neural recordings. *Nature Reviews Neuroscience*, 20(6):330–345.

Hoogenboom, N., Schoffelen, J.-M., Oostenveld, R., Parkes, L. M., and Fries, P. (2006). Localizing human visual gamma-band activity in frequency, time and space. *NeuroImage*, 29(3):764–773.

Horwitz, G. D., Chichilnisky, E. J., and Albright, T. D. (2004). Blue-yellow signals are enhanced by spatiotemporal luminance contrast in macaque V1. *Journal of Neurophysiology*, 93(4):2263–2278.

Horwitz, G. D. and Hass, C. A. (2012). Nonlinear analysis of macaque V1 color tuning reveals cardinal directions for cortical color processing. *Nature Neuroscience*, 15(6):913–919.

Hosoya, T., Baccus, S. A., and Meister, M. (2005). Dynamic predictive coding by the retina. *Nature*, 436(7047):71–77.

Huang, G., Ramachandran, S., Lee, T. S., and Olson, C. R. (2018). Neural correlate of visual familiarity in macaque area V2. *Journal of Neuroscience*, 38(42):8967–8975.

Hubel, D. H. and Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology*, 160(1):106–154.

Hubel, D. H. and Wiesel, T. N. (1965). Receptive fields and functional architecture in two nonstriate visual areas (18 and 19) of the cat. *Journal of Neurophysiology*, 28:229–89.

Huber, R., Mäki, H., Rosanova, M., Casarotto, S., Canali, P., Casali, A. G., Tononi, G., and Massimini, M. (2013). Human cortical excitability increases with time awake. *Cerebral Cortex*, 23(2):332–338.

Issa, E. B., Cadieu, C. F., and Dicarlo, J. J. (2018). Neural dynamics at successive stages of the ventral visual stream are consistent with hierarchical error signals. *eLife*, 7.

Jacobson, G. A., Rupprecht, P., and Friedrich, R. W. (2018). Experience-dependent plasticity of odor representations in the telencephalon of zebrafish. *Current Biology*, 28(1):1–14.e3.

Jadi, M. P. and Sejnowski, T. J. (2014). Cortical oscillations arise from contextual interactions that regulate sparse coding. *Proceedings of the National Academy of Sciences of the United States of America*, 111:6780–5.

James, T. W. and Gauthier, I. (2006). Repetition-induced changes in BOLD response reflect accumulation of neural activity. *Human Brain Mapping*, 27(1):37–46.

James, T. W., Humphrey, G. K., Gati, J. S., Menon, R. S., and Goodale, M. A. (2000). The effects of visual object priming on brain activation before and after recognition. *Current Biology*, 10(17):1017–1024.

James, W. (1890). The principles of psychology (Vols. 1 & 2). *New York Holt*, 118:688.

Jia, X., Smith, M. A., and Kohn, A. (2011). Stimulus selectivity and spatial coherence of gamma components of the local field potential. *Journal of Neuroscience*, 31(25):9390–9403.

Jia, X., Tanabe, S., and Kohn, A. (2013a). Gamma and the coordination of spiking activity in early visual cortex. *Neuron*, 77(4):762–774.

Jia, X., Xing, D., and Kohn, A. (2013b). No consistent relationship between gamma power and peak frequency in macaque primary visual cortex. *Journal of Neuroscience*, 33(1):17–25.

Johnson, E., Hawken, M., and Shapley, R. (2001). The spatial transformation of color in the primary visual cortex of the macaque monkey. *Nature Neuroscience*, 4(4):409–416.

Johnson, E. N., Hawken, M. J., and Shapley, R. (2008). The orientation selectivity of color-responsive neurons in macaque V1. *Journal of Neuroscience*, 28(32):8096–8106.

Jones, H. E., Andolina, I. M., Ahmed, B., Shipp, S. D., Clements, J. T. C., Grieve, K. L., Cudeiro, J., Salt, T. E., and Sillito, A. M. (2012). Differential feedback modulation of center and surround mechanisms in parvocellular cells in the visual thalamus. *Journal of Neuroscience*, 32(45):15946–15951.

Kaas, J. H. and Huerta, M. F. (1988). The subcortical visual system of primates. *Comparative Primate Biology*, 4:327–391.

Kaliukhovich, D. A., de Baene, W., and Vogels, R. (2013). Effect of adaptation on object representation accuracy in macaque inferior temporal cortex. *Journal of Cognitive Neuroscience*, 25(5):777–789.

Kaliukhovich, D. A. and Vogels, R. (2011). Stimulus repetition probability does not affect repetition suppression in macaque inferior temporal cortex. *Cerebral Cortex*, 21(7):1547–1558.

Kaliukhovich, D. A. and Vogels, R. (2012). Stimulus repetition affects both strength and synchrony of macaque inferior temporal cortical activity. *Journal of Neurophysiology*, 107(12):3509–3527.

Kaliukhovich, D. A. and Vogels, R. (2014). Neurons in macaque inferior temporal cortex show no surprise response to deviants in visual oddball sequences. *Journal of Neuroscience*, 34(38):12801–12815.

Kaliukhovich, D. A. and Vogels, R. (2016). Divisive normalization predicts adaptation-induced response changes in macaque inferior temporal cortex. *Journal of Neuroscience*, 36(22):6116–6128.

Kapadia, M., Westheimer, G., and Gilbert, C. (1999). Dynamics of spatial summation in primary visual cortex of alert monkeys. *Proceedings of the National Academy of Sciences*, 96(21):12073–12078.

Kapadia, M. K., Ito, M., Gilbert, C. D., and Westheimer, G. (1995). Improvement in visual sensitivity by changes in local context: Parallel studies in human observers and in V1 of alert monkeys. *Neuron*, 15(4):843–856.

Kaplan, E. S., Cooke, S. F., Komorowski, R. W., Chubykin, A. A., Thomazeau, A., Khibnik, L. A., Gavornik, J. P., and Bear, M. F. (2016). Contrasting roles for parvalbumin-expressing inhibitory neurons in two forms of adult visual cortical plasticity. *eLife*, 5.

Kar, K., Kubilius, J., Schmidt, K., Issa, E. B., and DiCarlo, J. J. (2019). Evidence that recurrent circuits are critical to the ventral stream's execution of core object recognition behavior. *Nature Neuroscience*, 22(6):974–983.

Kaschube, M., Schnabel, M., Löwel, S., Coppola, D. M., White, L. E., and Wolf, F. (2010). Universality in the evolution of orientation columns in the visual cortex. *Science*, 330(6007):1113–1116.

Kayser, C., Salazar, R. F., and Konig, P. (2003). Responses to natural scenes in cat V1. *Journal of Neurophysiology*, 90(3):1910–20.

Kim, H., Homann, J., Tank, D. W., and Berry, M. J. (2019). A long timescale stimulus history effect in the primary visual cortex. *bioRxiv*.

Kimura, F. (2009). Myelination and isochronicity in neural networks. *Frontiers in Neuroanatomy*, 3.

King, J. L., Lowe, M. P., Stover, K. R., Wong, A. A., and Crowder, N. A. (2016). Adaptive processes in thalamus and cortex revealed by silencing of primary visual cortex during contrast adaptation. *Current Biology*, 26(10):1295–1300.

Kissinger, S. T., Pak, A., Tang, Y., Masmanidis, S. C., and Chubykin, A. A. (2018). Oscillatory Encoding of Visual Stimulus Familiarity. *Journal of Neuroscience*, 38(27):6223–6240.

Knoblich, U., Siegle, J. H., Pritchett, D. L., and Moore, C. I. (2010). What do we gain from gamma? Local dynamic gain modulation drives enhanced efficacy and efficiency of signal transmission. *Frontiers in Human Neuroscience*, 4:185.

Koch, C., Rapp, M., and Segev, I. (1996). A brief history of time (constants). *Cerebral Cortex*, 6(2):93–101.

Kohn, A. (2007). Visual adaptation: Physiology, mechanisms, and functional benefits. *Journal of Neurophysiology*, 97(5):3155–3164.

Kohn, A., Coen-Cagli, R., Kanitscheider, I., and Pouget, A. (2016). Correlations and neuronal population information. *Annual Review of Neuroscience*, 39(1):237–256.

Kohn, A. and Movshon, J. A. (2003). Neuronal adaptation to visual motion in area MT of the macaque. *Neuron*, 39(4):681–691.

König, P., Engel, A. K., and Singer, W. (1996). Integrator or coincidence detector? The role of the cortical neuron revisited. *Trends in Neurosciences*, 19(4):130–137.

Kopell, N., Ermentrout, G. B., Whittington, M. A., and Traub, R. D. (2000). Gamma rhythms and beta rhythms have different synchronization properties. *Proceedings of the National Academy of Sciences*, 97:1867–1872.

Körding, K. P. and Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, 427(6971):244–7.

Korenbrot, J. I. (2012). Speed, sensitivity, and stability of the light response in rod and cone photoreceptors: Facts and models. *Progress in Retinal and Eye Research*, 31(5):442–466.

Korn, E. L., Troendle, J. F., McShane, L. M., and Simon, R. (2004). Controlling the number of false discoveries: application to high-dimensional genomic data. *Journal of Statistical Planning and Inference*, 124(2):379–398.

Korndörfer, C., Ullner, E., García-Ojalvo, J., and Pipa, G. (2017). Cortical spike synchrony as a measure of input familiarity. *Neural Computation*, 29(9):2491–2510.

Krauskopf, J., Williams, D. R., and Heeley, D. W. (1982). Cardinal directions of color space. *Vision Research*, 22(9):1123–1131.

Kruse, W. and Eckhorn, R. (1996). Inhibition of sustained gamma oscillations (35-80 Hz) by fast transient responses in cat visual cortex. *Proceedings of the National Academy of Sciences*, 93(12):6112–6117.

Kubilius, J., Schrimpf, M., Nayebi, A., Bear, D., Yamins, D. L. K., and DiCarlo, J. J. (2018). CORnet: Modeling the neural mechanisms of core object recognition. *bioRxiv*.

Kuhn, M., Wolf, E., Maier, J. G., Mainberger, F., Feige, B., Schmid, H., Bürklin, J., Maywald, S., Mall, V., Jung, N. H., Reis, J., Spiegelhalder, K., Klöppel, S., Sterr, A., Eckert, A., Riemann, D., Normann, C., and Nissen, C. (2016). Sleep recalibrates homeostatic and associative synaptic plasticity in the human cortex. *Nature Communications*, 7.

Kumar, S. and Hedges, S. B. (1998). A molecular timescale for vertebrate evolution. *Nature*, 392(6679):917–920.

Kuravi, P. and Vogels, R. (2017). Effect of adapter duration on repetition suppression in inferior temporal cortex. *Scientific Reports*, 7(1).

Lafer-Sousa, R., Liu, Y. O., Lafer-Sousa, L., Wiest, M. C., and Conway, B. R. (2012). Color tuning in alert macaque V1 assessed with fMRI and single-unit recording shows a bias toward daylight colors. *Journal of the Optical Society of America A*, 29(5):657.

Lamme, V. and Roelfsema, P. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences*, 23(11):571–579.

Lamme, V. A. (1995). The neurophysiology of figure-ground segregation in primary visual cortex. *Journal of Neuroscience*, 15(2):1605–1615.

Land, E. H. (1959). Color Vision and the Natural Image. Part I. *Proceedings of the National Academy of Sciences*, 45(4):115–129.

Landau, A. N. and Fries, P. (2012). Attention samples stimuli rhythmically. *Current Biology*, 22(11):1000–1004.

Laramée, M.-E. and Boire, D. (2015). Visual cortical areas of the mouse: comparison of parcellation and network structure with primates. *Frontiers in Neural Circuits*, 8.

Laurent, G. and Davidowitz, H. (1994). Encoding of olfactory information with oscillating neural assemblies. *Science*, 265(5180):1872–1875.

Laurent, G., MacLeod, K., Stopfer, M., and Wehr, M. (1999). Dynamic representation of odours by oscillating neural assemblies. In *Entomologia Experimentalis et Applicata*, volume 91, pages 7–18.

Laurent, G., Stopfer, M., Friedrich, R. W., Rabinovich, M. I., Volkovskii, A., and Abarbanel, H. D. (2001). Odor encoding as an active, dynamical process: experiments, computation, and theory. *Annual Review of Neuroscience*, 24:263–297.

Lazar, A., Lewis, C., Fries, P., Singer, W., and Nikolić, D. (2018). Visual exposure optimizes stimulus encoding in primary visual cortex. *bioRxiv*.

Lee, T. S. and Nguyen, M. (2001). Dynamics of subjective contour formation in the early visual cortex. *Proceedings of the National Academy of Sciences*, 98(4):1907–1911.

Legatt, A. D., Arezzo, J., and Vaughan, H. G. (1980). Averaged multiple unit activity as an estimate of phasic changes in local neuronal activity: effects of volume-conducted potentials. *Journal of Neuroscience Methods*, 2(2):203–217.

Leszczynski, M., Barczak, A., Kajikawa, Y., Ulbert, I., Falchier, A., Tal, I., Haegens, S., Melloni, L., Knight, R., and Schroeder, C. (2019). Dissociation of broadband high-frequency activity and neuronal firing in the neocortex. *bioRxiv*.

Letzkus, J. J., Wolff, S. B., Meyer, E. M., Tovote, P., Courtin, J., Herry, C., and Lüthi, A. (2011). A disinhibitory microcircuit for associative fear learning in the auditory cortex. *Nature*, 480(7377):331–335.

Levitan, I. B. and Kaczmarek, L. K. (2015). *The neuron: Cell and molecular biology, 4th Edition*. Oxford University Press.

Li, L., Miller, E. K., and Desimone, R. (1993). The representation of stimulus familiarity in anterior inferior temporal cortex. *Journal of Neurophysiology*, 69(6):1918–1929.

Li, W., Piëch, V., and Gilbert, C. D. (2004). Perceptual learning and top-down influences in primary visual cortex. *Nature Neuroscience*, 7(6):651–657.

Li, X., Chen, Y., Lashgari, R., Bereshpolova, Y., Swadlow, H. A., Lee, B. B., and Alonso, J. M. (2015). Mixing of chromatic and luminance retinal signals in primate area V1. *Cerebral Cortex*, 25(7):1920–1937.

Li, Z. (2002). A saliency map in primary visual cortex. *Trends in Cognitive Sciences*, 6(1):9–16.

Liang, H., Gong, X., Chen, M., Yan, Y., Li, W., and Gilbert, C. D. (2017). Interactions between feedback and lateral connections in the primary visual cortex. *Proceedings of the National Academy of Sciences*, 114(32):8637–8642.

Lichtman, J. W. and Denk, W. (2011). The big and the small: Challenges of imaging the brain's circuits. *Science*, 334(6056):618–623.

Lima, B., Singer, W., Chen, N. H., and Neuenschwander, S. (2010). Synchronization dynamics in response to plaid stimuli in monkey V1. *Cerebral Cortex*, 20:1556–1573.

Lima, B., Singer, W., and Neuenschwander, S. (2011). Gamma responses correlate with temporal expectation in monkey primary visual cortex. *Journal of Neuroscience*, 31(44):15919–15931.

Lisman, J. (2005). The theta/gamma discrete phase code occuring during the hippocampal phase precession may be a more general brain coding scheme. *Hippocampus*, 15(7):913–922.

Liu, X., Ward, B. D., Binder, J. R., Li, S. J., and Hudetz, A. G. (2014). Scale-free functional connectivity of the brain is maintained in anesthetized healthy participants but not in patients with unresponsive wakefulness syndrome. *PLoS One*, 9(3).

Livingstone, M. S. (1996). Oscillatory firing and interneuronal correlations in squirrel monkey striate cortex. *Journal of Neurophysiology*, 75:2467–2485.

Livingstone, S. and Hubel, H. (1984). Anatomy and physiology of a color system in the primate visual cortex. *Journal of Neuroscience*, 4(1):309–356.

Logothetis, N. K. and Wandell, B. A. (2004). Interpreting the BOLD Signal. *Annual Review of Physiology*, 66(1):735–769.

Lowet, E., Roberts, M., Hadjipapas, A., Peter, A., van der Eerden, J., and De Weerd, P. (2015). Input-Dependent Frequency Modulation of Cortical Gamma Oscillations Shapes Spatial Synchronization and Enables Phase Coding. *PLoS Computational Biology*, 11(2).

Lowet, E., Roberts, M. J., Bosman, C. A., Fries, P., and de Weerd, P. (2016). Areas V1 and V2 show microsaccade-related 3-4-Hz covariation in gamma power and frequency. *European Journal of Neuroscience*, 43(10):1286–1296.

Lowet, E., Roberts, M. J., Peter, A., Gips, B., and De Weerd, P. (2017). A quantitative theory of gamma synchronization in macaque V1. *eLife*, 6.

Lund, J. S., Angelucci, A., and Bressloff, P. C. (2003). Anatomical substrates for functional columns in macaque monkey primary visual cortex. *Cerebral Cortex*, 13(1):15–24.

Lund, J. S., Yoshioka, T., Lund, J. S., Yoshioka, T., Levitt, J. B., Lund, J. S., and Levitt, J. B. (1993). Comparison of intrinsic connectivity in different areas of macaque monkey cerebral cortex. *Cerebral Cortex*, 3(2):148–162.

Lundqvist, M., Herman, P., and Miller, E. K. (2018). Working memory: Delay activity, yes! Persistent activity? Maybe not. *Journal of Neuroscience*, 38(32):7013–7019.

Lundqvist, M., Rose, J., Herman, P., Brincat, S. L., Buschman, T. J., and Miller, E. K. (2016). Gamma and beta bursts underlie working memory. *Neuron*, 90(1):152–164.

Ma, W. J., Beck, J. M., Latham, P. E., and Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience*, 9(11):1432–1438.

Maass, W. (2016). Searching for principles of brain computation. *Current Opinion in Behavioral Sciences*, 11:81–92.

Maffei, L. and Fiorentini, A. (1976). The unresponsive regions of visual cortical receptive fields. *Vision Research*, 16(10).

Maier, A., Adams, G., Aura, C., and Leopold, D. (2010). Distinct superficial and deep laminar domains of activity in the visual cortex during rest and stimulation. *Frontiers in System Neuroscience*, 4:31.

Maier, A., Wilke, M., Aura, C., Zhu, C., Ye, F. Q., and Leopold, D. A. (2008). Divergence of fMRI and neural signals in V1 during perceptual suppression in the awake monkey. *Nature Neuroscience*, 11(10):1193–1200.

Malenka, R. C. and Bear, M. F. (2004). LTP and LTD: an embarrassment of riches. *Neuron*, 44(1):5–21.

Manahova, M. E., Spaak, E., and de Lange, F. P. (2019). Familiarity increases processing speed in the visual system. *bioRxiv*.

Mante, V. and Carandini, M. (2005). Mapping of stimulus energy in primary visual cortex. *Journal of Neurophysiology*, 94(1):788–798.

Mantini, D., Hasson, U., Betti, V., Perrucci, M. G., Romani, G. L., Corbetta, M., Orban, G. A., and Vanduffel, W. (2012). Interspecies activity correlations reveal functional correspondence between monkey and human brain areas. *Nature Methods*, 9(3):277–282.

Markov, N. T., Vezoli, J., Chameau, P., Falchier, A., Quilodran, R., Huissoud, C., Lamy, C., Misery, P., Giroud, P., Ullman, S., Barone, P., Dehay, C., Knoblauch, K., and Kennedy, H. (2014). Anatomy of hierarchy: Feedforward and feedback pathways in macaque visual cortex. *Journal of Comparative Neurology*, 522(1):225–259.

Markowitz, D. a., Collman, F., Brody, C. D., Hopfield, J. J., and Tank, D. W. (2008). Rate-specific synchrony: using noisy oscillations to detect equally active neurons. *Proceedings of the National Academy of Sciences of the United States of America*, 105(24):8422–7.

Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. Henry Holt and Co., Inc., New York, NY, USA.

Marr, D. and Brindley, G. S. (1970). A theory for cerebral neocortex. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 176(1043):161–234.

Martinez, L. M., Wang, Q., Reid, R. C., Pillai, C., Alonso, J. M., Sommer, F. T., and Hirsch, J. A. (2005). Receptive field structure varies with layer in the primary visual cortex. *Nature Neuroscience*, 8(3):372–379.

Masland, R. H. (2012). The neuronal organization of the retina. *Neuron*, 76(2):266–280.

Masquelier, T., Hugues, E., Deco, G., and Thorpe, S. J. (2009). Oscillations, phase-of-firing coding, and spike timing-dependent plasticity: an efficient learning scheme. *Journal of Neuroscience*, 29(43):13484–93.

McCormick, D. A. and Pape, H. C. (1990). Properties of a hyperpolarization activated cation current and its role in rhythmic oscillation in thalamic relay neurones. *Journal of Physiology*, 431(1):291–318.

McMahon, D. B. T. and Olson, C. R. (2007). Repetition suppression in monkey inferotemporal cortex: Relation to behavioral priming. *Journal of Neurophysiology*, 97(5):3532–3543.

McMahon, S. M. and Jackson, M. B. (2018). An inconvenient truth: Calcium sensors are calcium buffers. *Trends in Neurosciences*, 41(12):880–884.

Meister, M. and Berry, M. J. (1999). The neural code of the retina. *Neuron*, 22(3):435–50.

Melin, A. D., Chiou, K. L., Walco, E. R., Bergstrom, M. L., Kawamura, S., and Fedigan, L. M. (2017). Trichromacy increases fruit intake rates of wild capuchins (Cebus capucinus imitator). *Proceedings of the National Academy of Sciences*, 114(39):10402–10407.

Merker, B. (2013). Cortical gamma oscillations: The functional key is activation, not cognition. *Neuroscience and Biobehavioral Reviews*, 37(3):401–417.

Meyer, T. and Rust, N. (2018). Single-exposure visual memory judgments are reflected in inferotemporal cortex. *eLife*, 7.

Meyer, T., Walker, C., Cho, R. Y., and Olson, C. R. (2014). Image familiarization sharpens response dynamics of neurons in inferotemporal cortex. *Nature Neuroscience*, 17(10):1388–1394.

Michalareas, G., Vezoli, J., van Pelt, S., Schoffelen, J. M., Kennedy, H., and Fries, P. (2016). Alpha-beta and gamma rhythms subserve feedback and feedforward influences among human visual cortical areas. *Neuron*, 89(2):384–397.

Miller, E., Li, L., and Desimone, R. (1993). Activity of neurons in anterior inferior temporal cortex during a short-term memory task. *Journal of Neuroscience*, 13(4):1460–1478.

Miller, E. K. and Desimone, R. (1994). Parallel neuronal mechanisms for short-term memory. *Science*, 263(5146):520–522.

Miller, E. K., Lundqvist, M., and Bastos, A. M. (2018). Working Memory 2.0. *Neuron*, 100:463–475.

Miller, K., Zanos, S., Fetz, E., Den Nijs, M., and Ojemann, J. (2009a). Decoupling the cortical power spectrum reveals real-time representation of individual finger movements in humans. *Journal of Neuroscience*, 29(10):3132–3137.

Miller, K. J., Sorensen, L. B., Ojemann, J. G., and Den Nijs, M. (2009b). Power-law scaling in the brain surface electric potential. *PLoS Computational Biology*, 5(12).

Milner, P. M. (1974). A model for visual shape recognition. *Psychological Review*, 81(6):521.

Minsky, M. (1986). *The society of mind*. Simon and Schuster.

Mishkin, M., Ungerleider, L. G., and Macko, K. A. (1983). Object vision and spatial vision: two cortical pathways. *Trends in Neurosciences*, 6:414–417.

Mitra, P. P. and Pesaran, B. (1999). Analysis of dynamic brain imaging data. *Biophysical Journal*, 76:691–708.

Mitzdorf, U. (1985). Current source-density method and application in cat cerebral cortex: investigation of evoked potentials and EEG phenomena. *Physiological Reviews*, 65:37–100.

Mizokami, Y., Webster, S., and Webster, M. (2003). Seasonal variations in the color statistics of natural images. *Journal of Vision*, 3(9).

Mongillo, G., Barak, O., and Tsodyks, M. (2008). Synaptic theory of working memory. *Science*.

Moore, C. I., Carlen, M., Knoblich, U., and Cardin, J. A. (2010). Neocortical interneurons: from diversity, strength. *Cell*, 142:189–193.

Movshon, J. A. and Lennis, P. (1979). Pattern-selective adaptation in visual cortical neurones. *Nature*, 278(5707):850–852.

Movshon, J. A., Thompson, I. D., and Tolhurst, D. J. (1978). Receptive field organization of complex cells in the cat's striate cortex. *Journal of Physiology*, 283(1):79–99.

Müller, J. R., Metha, A. B., Krauskopf, J., and Lennie, P. (1999). Rapid adaptation in visual cortex to the structure of images. *Science*, 285(5432):1405–1408.

Muller, L., Chavane, F., Reynolds, J., and Sejnowski, T. J. (2018). Cortical travelling waves: Mechanisms and computational principles. *Nature Reviews Neuroscience*, 19(5):255–268.

Murphy, P. C., Duckett, S. G., and Sillito, A. M. (1999). Feedback connections to the lateral geniculate nucleus and cortical response properties. *Science*, 286(5444):1552–1554.

Murty, D. V., Shirhatti, V., Ravishankar, P., and Ray, S. (2018). Large visual stimuli induce two distinct gamma oscillations in primate visual cortex. *Journal of Neuroscience*, 38(11):2730–2744.

Muthukumaraswamy, S. D. and Singh, K. D. (2013). Visual gamma oscillations: The effects of stimulus type, visual field coverage and stimulus motion on MEG and EEG recordings. *NeuroImage*, 69:223–230.

Natan, R. G., Rao, W., and Geffen, M. N. (2017). Cortical interneurons differentially shape frequency tuning following adaptation. *Cell Reports*, 21(4):878–890.

Neuenschwander, S. and Singer, W. (1996). Long-range synchronization of oscillatory light responses in the cat retina and lateral geniculate nucleus. *Nature*, 379(6567):728–733.

Nienborg, H. (2004). Receptive field size in V1 neurons limits acuity for perceiving disparity modulation. *Journal of Neuroscience*, 24(9):2065–2076.

Niessing, M., Galuske, R. A., Niessing, J., Ebisch, B., Schmidt, K. E., and Singer, W. (2005). Hemodynamic signals correlate tightly with synchronized gamma oscillations. *Science*, 309(5736):948–51.

Nijhawan, R. (2008). Visual prediction: Psychophysics and neurophysiology of compensation for time delays. *Behavioral and Brain Sciences*, 31(2).

Nikolić, D. (2009). Model this! Seven empirical phenomena missing in the models of cortical oscillatory dynamics. In *Proceedings of the International Joint Conference on Neural Networks*, pages 2272–2279.

Nikolić, D., Usler, S. H., Singer, W., and Maass, W. (2009). Distributed fading memory for stimulus properties in the primary visual cortex. *PLoS Biology*, 7(12).

Nir, Y., Fisch, L., Mukamel, R., Gelbard-Sagiv, H., Arieli, A., Fried, I., and Malach, R. (2007). Coupling between neuronal firing rate, gamma LFP, and BOLD fMRI is related to interneuronal correlations. *Current Biology*, 17(15):1275–1285.

Nowak, L. G., Azouz, R., Sanchez-Vives, M. V., Gray, C. M., and McCormick, D. A. (2003). Electrophysiological classes of cat primary visual cortical neurons in vivo as revealed by quantitative analyses. *Journal of Neurophysiology*, 89:1541–1566.

Nurminen, L. and Angelucci, A. (2014). Multiple components of surround modulation in primary visual cortex: Multiple neural circuits with multiple functions? *Vision Research*, 104:47–56.

Olsen, S. R. and Wilson, R. I. (2008). Lateral presynaptic inhibition mediates gain control in an olfactory circuit. *Nature*, 452(7190):956–60.

Olshausen, B. A. and Field, D. J. (2005). How close are we to understanding V1? *Neural Computation*, 17(8):1665–1699.

Onorato, I., Neuenschwander, S., Hoy, J., Lima, B., Rocha, K.-S., Broggini, A. C., Uran, C., Spyropoulos, G., Womelsdorf, T., Fries, P., Niell, C., Singer, W., and Vinck, M. (2019). A distinct class of bursting neurons with strong gamma synchronization and stimulus selectivity in monkey V1. *bioRxiv*.

Oostenveld, R., Fries, P., Maris, E., Schoffelen, J.-M., Oostenveld, R., Fries, P., Maris, E., and Schoffelen, J.-M. (2011). FieldTrip: Open source software for advanced analysis of meg, eeg, and invasive electrophysiological data. *Computational Intelligence and Neuroscience*.

Orban, G. A., Van Essen, D., and Vanduffel, W. (2004). Comparative mapping of higher visual areas in monkeys and humans. *Trends in Cognitive Sciences*, 8(7):315–324.

Orekhova, E. V., Butorina, A. V., Sysoeva, O. V., Prokofyev, A. O., Nikolaeva, A. Y., and Stroganova, T. A. (2015). Frequency of gamma oscillations in humans is modulated by velocity of visual motion. *Journal of Neurophysiology*, 114(1):244–255.

Osorio, D. and Vorobyev, M. (1996). Colour vision as an adaptation to frugivory in primates. *Proceedings of the Royal Society B: Biological Sciences*, 263(1370):593–599.

Palmer, S. E., Marre, O., Berry, M. J., and Bialek, W. (2015). Predictive information in a sensory population. *Proceedings of the National Academy of Sciences*, 112(22):6908–6913.

Palmigiano, A., Geisel, T., Wolf, F., and Battaglia, D. (2017). Flexible information routing by transient synchrony. *Nature Neuroscience*, 20(7):1014–1022.

Patterson, C. A., Wissig, S. C., and Kohn, A. (2013). Distinct effects of brief and prolonged adaptation on orientation tuning in primary visual cortex. *Journal of Neuroscience*, 33(2):532–543.

Patterson, C. A., Wissig, S. C., and Kohn, A. (2014). Adaptation disrupts motion integration in the primate dorsal stream. *Neuron*, 81(3):674–686.

Paxinos, G., Huang, X., Petrides, M., and Toga, A. (2009). *The rhesus monkey brain in stereotaxic coordinates*. Academic Press.

Peissig, J. J., Singer, J., Kawasaki, K., and Sheinberg, D. L. (2007). Effects of long-term object familiarity on event-related potentials in the monkey. *Cerebral Cortex*, 17(6):1323–1334.

Peli, E. (1990). Contrast in complex images. *Journal of the Optical Society of America A*, 7(10):2032.

Penttonen, M. and Buzsáki, G. (2003). Natural logarithmic relationship between brain oscillators. *Thalamus and Related Systems*, 2(2):145–152.

Perrenoud, Q., Pennartz, C. M., and Gentet, L. J. (2016). Membrane potential dynamics of spontaneous and visually evoked gamma activity in V1 of awake mice. *PLoS Biology*, 14(2).

Perry, G., Hamandi, K., Brindley, L. M., Muthukumaraswamy, S. D., and Singh, K. D. (2013). The properties of induced gamma oscillations in human visual cortex show individual variability in their dependence on stimulus size. *NeuroImage*, 68:83–92.

Pesaran, B., Vinck, M., Einevoll, G., Sirota, A., Fries, P., Siegel, M., Truccolo, W., Schroeder, C., and Srinivasan, R. (2018). Investigating large-scale brain dynamics using field potential recordings: analysis and interpretation. *Nature Neuroscience*.

Peter, A., Uran, C., Klon-Lipok, J., Roese, R., van Stijn, S., Barnes, W., Dowdall, J. R., Singer, W., Fries, P., and Vinck, M. (2019). Surface color and predictability determine contextual modulation of V1 firing and gamma oscillations. *eLife*, 8.

Petry, H. M. and Bickford, M. E. (2019). The second visual system of the tree shrew. *Journal of Comparative Neurology*, 527(3):679–693.

Poort, J., Raudies, F., Wannig, A., Lamme, V. A., Neumann, H., and Roelfsema, P. R. (2012). The role of attention in figure-ground segregation in areas V1 and V4 of the visual cortex. *Neuron*, 75(1):143–156.

Priebe, N. J. and Ferster, D. (2012). Mechanisms of neuronal computation in mammalian visual cortex. *Neuron*, 75(2):194–208.

Priebe, N. J. and Lisberger, S. G. (2002). Constraints on the source of short-term motion adaptation in macaque area MT. II. Tuning of neural circuit mechanisms. *Journal of Neurophysiology*, 88(1):370–382.

Pritchard, W. S. (1992). The brain in fractal time: 1/f-like power spectrum scaling of the human electroencephalogram. *International Journal of Neuroscience*, 66(1-2):119–129.

Pryluk, R., Kfir, Y., Gelbard-Sagiv, H., Fried, I., and Paz, R. (2019). A tradeoff in the neural code across regions and species. *Cell*, 176(3):597–609.

Quax, S. and Van Gerven, M. (2018). Emergent mechanisms of evidence integration in recurrent neural networks. *PLoS One*, 13(10).

Rainer, G. and Miller, E. K. (2000). Effects of visual experience on the representation of objects in the prefrontal cortex. *Neuron*, 27(1):179–189.

Rakic, P. (1994). Corticogenesis in human and nonhuman primates. In *The Cognitive Neurosciences*, pages 127–145. MIT press.

Rao, R. P. and Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1):79–87.

Ray, S. and Maunsell, J. H. (2011). Different origins of gamma rhythm and high-gamma activity in macaque visual cortex. *PLoS Biology*, 9(4):e1000610.

Ray, S. and Maunsell, J. H. R. (2010). Differences in gamma frequencies across visual cortex restrict their possible use in computation. *Neuron*, 67:885–96.

Re, D., Inbar, M., Richter, C. G., and Landau, A. N. (2019). Feature-based attention samples stimuli rhythmically. *Current Biology*, 29(4):693–699.e4.

Reig, R., Gallego, R., Nowak, L. G., and Sanchez-Vives, M. V. (2006). Impact of cortical network activity on short-term synaptic depression. *Cerebral Cortex*, 16(5):688–695.

Richter, C. G., Coppola, R., and Bressler, S. L. (2018). Top-down beta oscillatory signaling conveys behavioral context in early visual cortex. *Scientific Reports*, 8(1):6991.

Ringach, D. and Shapley, R. (2004). Reverse correlation in neurophysiology. *Cognitive Science*, 28(2):147–166.

Ringo, J. L. (1996). Stimulus specific adaptation in inferior temporal and medial temporal cortex of the monkey. *Behavioural Brain Research*, 76(1-2):191–197.

Roberts, M., Delicato, L. S., Herrero, J., Gieselmann, M. A., and Thiele, A. (2007). Attention alters spatial integration in macaque V1 in an eccentricity-dependent manner. *Nature Neuroscience*, 10(11):1483–1491.

Roberts, M. J., Lowet, E., Brunet, N. M., Ter Wal, M., Tiesinga, P., Fries, P., and De Weerd, P. (2013). Robust gamma coherence between macaque V1 and V2 by dynamic frequency matching. *Neuron*, 78(3):523–536.

Rockland, K. S. and Pandya, D. N. (1979). Laminar origins and terminations of cortical connections of the occipital lobe in the rhesus monkey. *Brain Research*, 179(1):3–20.

Roe, A. W., Chelazzi, L., Connor, C. E., Conway, B. R., Fujita, I., Gallant, J. L., Lu, H., and Vanduffel, W. (2012). Toward a unified theory of visual area V4. *Neuron*, 74(1):12–29.

Roelfsema, P. R., Lamme, V. A., and Spekreijse, H. (2004). Synchrony and covariation of firing rates in the primary visual cortex during contour grouping. *Nature Neuroscience*, 7(9):982–991.

Roelfsema, P. R., van Ooyen, A., and Watanabe, T. (2010). Perceptual learning rules based on reinforcers and attention. *Trends in Cognitive Sciences*, 14(2):64–71.

Rohenkohl, G., Bosman, C. A., and Fries, P. (2018). Gamma synchronization between V1 and V4 improves behavioral performance. *Neuron*, 100(4):953–963.e3.

Rolf, K. (2011). Scalable Brain Atlas: From stereotaxic coordinate to delineated brain region. *Frontiers in Neuroscience*, 4.

Rolls, E. T., Webb, T. J., and Deco, G. (2012). Communication before coherence. *European Journal of Neuroscience*, 36(5):2689–2709.

Rols, G., Tallon-Baudry, C., Girard, P., Bertrand, O., and Bullier, J. (2001). Cortical mapping of gamma oscillations in areas V1 and V4 of the macaque monkey. *Visual Neuroscience*, 18:527–540.

Rothman, J. S., Cathala, L., Steuber, V., and Silver, R. A. (2009). Synaptic depression enables neuronal gain control. *Nature*, 457(7232):1015–1018.

Rubehn, B., Bosman, C., Oostenveld, R., Fries, P., and Stieglitz, T. (2009). A MEMS-based flexible multichannel ECoG-electrode array. *Journal of Neural Engineering*, 6:036003.

Rust, N. C. and Movshon, J. A. (2005). In praise of artifice. *Nature Neuroscience*, 8(12):1647–1650.

Salami, M., Itami, C., Tsumoto, T., and Kimura, F. (2003). Change of conduction velocity by regional myelination yields constant latency irrespective of distance between thalamus and cortex. *Proceedings of the National Academy of Sciences*, 100(10):6174–6179.

Salin, P. A. and Bullier, J. (1995). Corticocortical connections in the visual system: structure and function. *Physiological Reviews*, 75(1):107–154.

Salinas, E. and Sejnowski, T. J. (2001). Correlated neuronal activity and the flow of neural information. *Nature Reviews Neuroscience*, 2:539–50.

Sanchez-Vives, M., Nowak, L., and McCormick, D. (2000a). Cellular mechanisms of long-lasting adaptation in visual cortical neurons in vitro. *Journal of Neuroscience*, 20(11):4286–4299.

Sanchez-Vives, M. V. and McCormick, D. A. (2000). Cellular and network mechanisms of rhytmic recurrent activity in neocortex. *Nature Neuroscience*, 3(10):1027–1034.

Sanchez-Vives, M. V., Nowak, L. G., and McCormick, D. A. (2000b). Membrane mechanisms underlying contrast adaptation in cat area 17 in vivo. *Journal of Neuroscience*, 20(11):4267–4285.

Santos, L. R., Hauser, M. D., and Spelke, E. S. (2001). Recognition and categorization of biologically significant objects by rhesus monkeys (Macaca mulatta): The domain of food. *Cognition*, 82(2):127–155.

Sawamura, H., Georgieva, S., Vogels, R., Vanduffel, W., and Orban, G. A. (2005). Using functional magnetic resonance imaging to assess adaptation and size invariance of shape processing by humans and monkeys. *Journal of Neuroscience*, 25(17):4294–4306.

Sawamura, H., Orban, G. A., and Vogels, R. (2006). Selectivity of neuronal adaptation does not match response selectivity: A single-cell study of the fMRI adaptation paradigm. *Neuron*, 49(2):307–318.

Scelfo, B. and Buffelli, M. R. (2009). Developmental axonal pruning and synaptic plasticity. In *The Sticky Synapse: Cell Adhesion Molecules and Their Role in Synapse Formation and Maintenance*, pages 107–140. Springer.

Scheeringa, R., Koopmans, P. J., van Mourik, T., Jensen, O., and Norris, D. G. (2016). The relationship between oscillatory EEG activity and the laminar-specific BOLD signal. *Proceedings of the National Academy of Sciences*, 113(24):6761–6766.

Scheffer-Teixeira, R. and Tort, A. B. (2016). On cross-frequency phase-phase coupling between theta and gamma oscillations in the hippocampus. *eLife*, 5:e20515.

Schluppeck, D. and Engel, S. A. (2002). Color opponent neurons in V1: a review and model reconciling results from imaging and single-unit recording. *Journal of Vision*, 2(6):5.

Schmid, M. C., Schmiedt, J. T., Peters, A. J., Saunders, R. C., Maier, A., and Leopold, D. A. (2013). Motion-sensitive responses in visual area V4 in the absence of primary visual cortex. *Journal of Neuroscience*, 33(48):18740–18745.

Schoffelen, J. M., Oostenveld, R., and Fries, P. (2005). Neuronal coherence as a mechanism of effective corticospinal interaction. *Science*, 308(5718):111–113.

Schoups, A., Vogels, R., Qian, N., and Orban, G. (2001). Practising orientation identification improves orientation coding in V1 neurons. *Nature*, 412(6846):549–553.

Schwartz, O., Hsu, A., and Dayan, P. (2007). Space and time in visual context. *Nature Reviews Neuroscience*, 8(7):522–535.

Schwartz, O. and Simoncelli, E. P. (2001). Natural signal statistics and sensory gain control. *Nature Neuroscience*, 4(8):819–25.

Schwartz, O., Snow, M., and Coen-Cagli, R. (2017). Adaptation in the visual cortex: A case for probing neuronal populations with natural stimuli. *F1000Research*, 6.

Schwiedrzik, C. M. and Freiwald, W. A. (2017). High-level prediction signals in a low-level area of the macaque face-processing hierarchy. *Neuron*, 96(1):89–97.e4.

Sejnowski, T. J. and Paulsen, O. (2006). Network oscillations: emerging computational principles. *Journal of Neuroscience*, 26:1673–1676.

Self, M. W., van Kerkoerle, T., Supèr, H., and Roelfsema, P. R. (2013). Distinct roles of the cortical layers of area V1 in figure-ground segregation. *Current Biology*, 23(21):2121–2129.

Serre, T., Kouh, M., Cadieu, C., Knoblich, U., Kreiman, G., and Poggio, T. (2005). A theory of object recognition: Computations and circuits in the feedforward path of the ventral stream in primate visual cortex. *Artificial Intelligence*, pages 1–130.

Shadlen, M. N. and Movshon, J. A. (1999). Synchrony unbound: a critical evaluation of the temporal binding hypothesis. *Neuron*, 24(1):67–77, 111–25.

Shapley, R. (2019). Physiology of color vision in primates. In *Oxford Research Encyclopedia of Neuroscience. Online Publication.*

Shapley, R. and Hawken, M. (2002). Neural mechanisms for color perception in the primary visual cortex. *Current Opinion in Neurobiology*, 12(4):426–432.

Shapley, R. and Hawken, M. J. (2011). Color in the cortex: single-and double-opponent cells. *Vision Research*, 51(7):701–717.

Sharpee, T. O., Sugihara, H., Kurgansky, A. V., Rebrik, S. P., Stryker, M. P., and Miller, K. D. (2006). Adaptive filtering enhances information transmission in visual cortex. *Nature*, 439(7079):936–942.

Shepherd, S. V., Steckenfinger, S. A., Hasson, U., and Ghazanfar, A. A. (2010). Human-Monkey gaze correlations reveal convergent and divergent patterns of movie viewing. *Current Biology*, 20(7):649–656.

Shirhatti, V. and Ray, S. (2018). Long-wavelength (reddish) hues induce unusually large gamma oscillations in the primate primary visual cortex. *Proceedings of the National Academy of Sciences*, 115(17):4489–4494.

Sillito, A. M., Grieve, K. L., Jones, H. E., Cudelro, J., and Davis, J. (1995). Visual cortical mechanisms detecting focal orientation discontinuities. *Nature*, 378.

Sillito, A. M. and Jones, H. E. (2002). Corticothalamic interactions in the transfer of visual information. In *Philosophical Transactions of the Royal Society B: Biological Sciences*, volume 357, pages 1739–1752.

Simoncelli, E. P. and Olshausen, B. A. (2001). Natural image statistics and neural representation. *Annual Review of Neuroscience*, 24:1193–216.

Singer, W. (1999). Neuronal synchrony: a versatile code for the definition of relations? *Neuron*, 24:49–65.

Singer, W. (2008). Synchronous oscillations and memory formation. In *Learning and Memory: A Comprehensive Reference*, pages 721–728. Elsevier.

Singer, W. (2018). Neuronal oscillations: unavoidable and useful? *European Journal of Neuroscience*, 48(7):2389–2398.

Singer, W. and Gray, C. M. (1995). Visual feature integration and the temporal correlation hypothesis. *Annual Review of Neuroscience*, 18:555–586.

Smith, A. (2002). Estimating receptive field size from fmri data in human striate and extrastriate visual cortex. *Cerebral Cortex*, 11(12):1182–1190.

Snow, M., Coen-Cagli, R., and Schwartz, O. (2016). Specificity and timescales of cortical adaptation as inferences about natural movie statistics. *Journal of Vision*, 16(13).

Softky, W. (1994). Sub-millisecond coincidence detection in active dendritic trees. *Neuroscience*, 58:13–41.

Sohal, V. S., Zhang, F., Yizhar, O., and Deisseroth, K. (2009). Parvalbumin neurons and gamma rhythms enhance cortical circuit performance. *Nature*, 459:698–702.

Solomon, S. G. (2006). Suppressive surrounds and contrast gain in magnocellular-pathway retinal ganglion cells of macaque. *Journal of Neuroscience*, 26(34):8715–8726.

Solomon, S. G. and Kohn, A. (2014). Moving sensory adaptation beyond suppressive effects in single neurons. *Current Biology*, 24(20):R1012–R1022.

Solomon, S. G. and Lennie, P. (2007). The machinery of colour vision. *Nature Reviews Neuroscience*, 8(4):276–286.

Solomon, S. G., Peirce, J. W., Dhruv, N. T., and Lennie, P. (2004). Profound contrast adaptation early in the visual pathway. *Neuron*, 42(1):155–62.

Sparks, D., Rohrer, W. H., and Zhang, Y. (2000). The role of the superior colliculus in saccade initiation: A study of express saccades and the gap effect. *Vision Research*, 40(20):2763–2777.

Spillmann, L., Dresp-Langley, B., and Tseng, C.-H. (2015). Beyond the classical receptive field: The effect of contextual stimuli. *Journal of Vision*, 15(9):7.

Spratling, M. W. (2010). Predictive coding as a model of response properties in cortical area V1. *Journal of Neuroscience*, 30(9):3531–3543.

Spyropoulos, G., Dowdall, J. R., Schölvinck, M. L., Bosman, C. A., Lima, B., Peter, A., Onorato, I., Klon-Lipok, J., Roese, R., Neuenschwander, S., Singer, W., Vinck, M., and Fries, P. (2019). On the relationship between gamma amplitude and frequency in awake macaque V1. *submitted*.

Srinivasan, S., Carlo, C. N., and Stevens, C. F. (2015). Predicting visual acuity from the structure of visual cortex. *Proceedings of the National Academy of Sciences*, 112(25):7815–7820.

Standing, L. (1973). Learning 10,000 pictures. *Quarterly Journal of Experimental Psychology*, 25(2):207–222.

Stoelzel, C. R., Huff, J. M., Bereshpolova, Y., Zhuang, J., Hei, X., Alonso, J.-M., and Swadlow, H. A. (2015). Hour-long adaptation in the awake early visual system. *Journal of Neurophysiology*, 114(2):1172–1182.

Stopfer, M., Bhagavan, S., Smith, B. H., and Laurent, G. (1997). Impaired odour discrimination on desynchronization of odour-encoding neural assemblies. *Nature*, 390(6655):70–74.

Stopfer, M. and Laurent, G. (1999). Short-term memory in olfactory network dynamics. *Nature*, 402(6762):664–668.

Stosiek, C., Garaschuk, O., Holthoff, K., and Konnerth, A. (2003). In vivo two-photon calcium imaging of neuronal networks. *Proceedings of the National Academy of Sciences*, 100(12):7319–7324.

Stringer, C., Pachitariu, M., Steinmetz, N., Carandini, M., and Harris, K. D. (2019). High-dimensional geometry of population responses in visual cortex. *Nature*, 571:361–365.

Sugase-Miyamoto, Y., Liu, Z., Wiener, M. C., Optican, L. M., and Richmond, B. J. (2008). Short-term memory trace in rapidly adapting synapses of inferior temporal cortex. *PLoS Computational Biology*, 4(5).

Sun, H.-J., Zhao, J., Southall, T. L., and Xu, B. (2002). Contextual influences on the directional responses of tectal cells in pigeons. *Visual Neuroscience*, 19(2):133–144.

Supèr, H., Spekreijse, H., and Lamme, V. A. (2003). Figure-ground activity in primary visual cortex (V1) of the monkey matches the speed of behavioral response. *Neuroscience Letters*, 344(2):75–78.

Surridge, A. K., Osorio, D., and Mundy, N. I. (2003). Evolution and selection of trichromatic vision in primates. *Trends in Ecology and Evolution*, 18(4):198–205.

Sutter, M. L., Schreiner, C. E., McLean, M., O'connor, K. N., and Loftus, W. C. (1999). Organization of inhibitory frequency receptive fields in cat primary auditory cortex. *Journal of Neurophysiology*, 82(5):2358–2371.

Tailby, C., Solomon, S. G., Dhruv, N. T., and Lennie, P. (2008a). Habituation reveals fundamental chromatic mechanisms in striate cortex of macaque. *Journal of Neuroscience*, 28(5):1131–1139.

Tailby, C., Solomon, S. G., and Lennie, P. (2008b). Functional asymmetries in visual pathways carrying s-cone signals in macaque. *Journal of Neuroscience*, 28(15):4078–4087.

Thomas, E. and French, R. (2017). Grandmother cells: much ado about nothing. *Language, Cognition and Neuroscience*, 32(3):342–349.

Thomsen, K., Offenhauser, N., and Lauritzen, M. (2004). Principal neuron spiking: Neither necessary nor sufficient for cerebral blood flow in rat cerebellum. *Journal of Physiology*, 560(1):181–189.

Tiesinga, P. and Sejnowski, T. J. (2009). Cortical enlightenment: are attentional gamma oscillations driven by ING or PING? *Neuron*, 63:727–732.

Todorovic, A., van Ede, F., Maris, E., and de Lange, F. P. (2011). Prior Expectation Mediates Neural Adaptation to Repeated Sounds in the Auditory Cortex: An MEG Study. *Journal of Neuroscience*, 31(25):9118–9123.

Tononi, G. (2009). Slow wave homeostasis and synaptic plasticity. *Journal of Clinical Sleep Medicine*, 5.

Tononi, G., Sporns, O., and Edelman, G. M. (2002). Measures of degeneracy and redundancy in biological networks. *Proceedings of the National Academy of Sciences*, 96(6):3257–3262.

Tootell, R. B. H., Tsao, D., and Vanduffel, W. (2003). Neuroimaging weighs in: humans meet macaques in "primate" visual cortex. *Journal of Neuroscience*, 23(10):3981–9.

Traub, R. D., Spruston, N., Soltesz, I., Konnerth, A., Whittington, M. A., and Jefferys, J. G. (1998). Gamma-frequency oscillations: A neuronal population phenomenon, regulated by synaptic and intrinsic cellular processes, and inducing synaptic plasticity. *Progress in Neurobiology*, 55(6):563–575.

Treisman, A. (1996). The binding problem. *Current Opinion in Neurobiology*, 6:171–78.

Trott, A. R. and Born, R. T. (2015). Input-gain control produces feature-specific surround suppression. *Journal of Neuroscience*, 35(12):4973–4982.

Ts'o, D. Y. and Gilbert, C. (1988). The organization of chromatic and spatial interactions in the primate striate cortex. *Journal of Neuroscience*, 8(5):1712–1727.

Turner, M. H., Sanchez Giraldo, L. G., Schwartz, O., and Rieke, F. (2019). Stimulus- and goal-oriented frameworks for understanding natural vision. *Nature Neuroscience*, 22(1):15–24.

Valiant, L. G. (2014). What must a global theory of cortex explain? *Current Opinion in Neurobiology*, 25:15–19.

van Den Bergh, G., Zhang, B., Arckens, L., and Chino, Y. M. (2010). Receptive-field properties of V1 and V2 neurons in mice and macaque monkeys. *Journal of Comparative Neurology*, 518(11):2051–2070.

van Essen, D. C. (2004). Organization of Visual Areas in Macaque and Human Cerebral Cortex. *Cortex*, 1:1–19.

van Essen, D. C., Anderson, C. H., and Felleman, D. J. (1992). Information processing in the primate visual system: An integrated systems perspective. *Science*, 255(5043):419–423.

van Essen, D. C. and Maunsell, J. H. (1983). Hierarchical organization and functional streams in the visual cortex. *Trends in Neurosciences*, 6(C):370–375.

van Hooser, S. D. (2007). Similarity and diversity in visual cortex: Is there a unifying theory of cortical computation? *Neuroscientist*, 13(6):639–656.

van Kerkoerle, T., Self, M. W., Dagnino, B., Gariel-Mathis, M.-A., Poort, J., van der Togt, C., and 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.

van Pelt, S., Boomsma, D. I., and Fries, P. (2012). Magnetoencephalography in twins reveals a strong genetic determination of the peak frequency of visually induced gamma-band synchronization. *Journal of Neuroscience*, 32(10):3388–3392.

van Pelt, S. and Fries, P. (2013). Visual stimulus eccentricity affects human gamma peak frequency. *NeuroImage*, 78:439–447.

van Wezel, R. J. A. and Britten, K. H. (2002). Motion Adaptation in Area MT. *Journal of Neurophysiology*, 88(6):3469–3476.

van Wingerden, M., Vinck, M., Lankelma, J., and Pennartz, C. (2010). Learning-associated gamma-band phase-locking of action-outcome selective neurons in orbitofrontal cortex. *Journal of Neuroscience*, 30:10025–10038.

van Wingerden, M., Vinck, M., Tijms, V., Ferreira, I. R., Jonker, A. J., and Pennartz, C. M. (2012). NMDA Receptors Control Cue-Outcome Selectivity and Plasticity of Orbitofrontal Firing Patterns during Associative Stimulus-Reward Learning. *Neuron*, 76(4):813–825.

Varela, F., Lachaux, J. P., Rodriguez, E., and Martinerie, J. (2001). The brainweb: phase synchronization and large-scale integration. *Nature Reviews Neuroscience*, 2:229–239.

Vázquez, P., Cano, M., and Acuña, C. (2017). Discrimination of line orientation in humans and monkeys. *Journal of Neurophysiology*, 83(5):2639–2648.

Vega-Bermudez, F. and Johnson, K. O. (1999). Surround suppression in the responses of primate SA1 and RA mechanoreceptive afferents mapped with a probe array. *Journal of Neurophysiology*, 81(6):2711–2719.

Veilleux, C. C. and Kirk, E. C. (2014). Visual acuity in mammals: Effects of eye size and ecology. *Brain, Behavior and Evolution*, 83(1):43–53.

Veit, J., Hakim, R., Jadi, M. P., Sejnowski, T. J., and Adesnik, H. (2017). Cortical gamma band synchronization through somatostatin interneurons. *Nature Neuroscience*, 20(7):951.

Vinck, M., Batista-Brito, R., Knoblich, U., and Cardin, J. A. (2015). Arousal and locomotion make distinct contributions to cortical activity patterns and visual encoding. *Neuron*, 86(3):740–754.

Vinck, M., Battaglia, F. P., Womelsdorf, T., and Pennartz, C. (2012). Improved measures of phase-coupling between spikes and the Local Field Potential. *Journal of Computational Neuroscience*, 33(1):53–75.

Vinck, M. and Bosman, C. A. (2016). More gamma more predictions: Gamma-synchronization as a key mechanism for efficient integration of classical receptive field inputs with surround predictions. *Frontiers in Systems Neuroscience*, 10:35.

Vinck, M., Lima, B., Womelsdorf, T., Oostenveld, R., Singer, W., Neuenschwander, S., and Fries, P. (2010a). Gamma-phase shifting in awake monkey visual cortex. *Journal of Neuroscience*, 30:1250–1257.

Vinck, M., van Wingerden, M., Womelsdorf, T., Fries, P., and Pennartz, C. M. (2010b). The pairwise phase consistency: a bias-free measure of rhythmic neuronal synchronization. *Neuroimage*, 51:112–122.

Vinck, M., Womelsdorf, T., Buffalo, E. A., Desimone, R., and Fries, P. (2013a). Attentional modulation of cell-class-specific gamma-band synchronization in awake monkey area V4. *Neuron*, 80(4):1077–1089.

Vinck, M., Womelsdorf, T., and Fries, P. (2013b). Gamma-band synchronization and information transmission. In Quiroga-Quian, R. and Panzeri, S., editors, *Principles of Neural Coding*. CRC Press.

Vinje, W. E. and Gallant, J. L. (2000). Sparse coding and decorrelation in primary visual cortex during natural vision. *Science*, 287(5456):1273–1276.

Vinken, K., Vogels, R., and Op de Beeck, H. (2017). Recent visual experience shapes visual processing in rats through stimulus-specific adaptation and response enhancement. *Current Biology*, 27(6):914–919.

Viswanathan, A. and Freeman, R. D. (2007). Neurometabolic coupling in cerebral cortex reflects synaptic more than spiking activity. *Nature Neuroscience*, 10(10):1308–1312.

Vogels, R. (2016). Sources of adaptation of inferior temporal cortical responses. *Cortex*, 80:185–195.

Vogels, R. and Orban, G. A. (1990). How well do response changes of striate neurons signal differences in orientation: a study in the discriminating monkey. *Journal of Neuroscience*, 10(11):3543–58.

Vogels, R. and Orban, G. A. (1994). Activity of inferior temporal neurons during orientation discrimination with successively presented gratings. *Journal of Neurophysiology*, 71(4):1428–1451.

von der Heydt, R., Peterhans, E., and Baumgartner, G. (1984). Illusory contours and cortical neuron responses. *Science*, 224(4654):1260–1262.

von der Malsburg, C. (1994). *The correlation theory of brain function*. Springer.

von Stein, A. and Sarnthein, J. (2000). Different frequencies for different scales of cortical integration: from local gamma to long range alpha/theta synchronization. *International Journal of Psychophysiology*, 38(3):301–313.

Vyazovskiy, V. V., Cirelli, C., Pfister-Genskow, M., Faraguna, U., and Tononi, G. (2008). Molecular and electrophysiological evidence for net synaptic potentiation in wake and depression in sleep. *Nature Neuroscience*, 11(2):200–208.

Wachtler, T., Sejnowski, T. J., and Albright, T. D. (2003). Representation of color stimuli in awake macaque primary visual cortex. *Neuron*, 37(4):681–691.

Wacongne, C., Changeux, J.-P., and Dehaene, S. (2012). A neuronal model of predictive coding accounting for the mismatch negativity. *Journal of Neuroscience*, 32(11):3665–3678.

Waitt, C., Gerald, M. S., Little, A. C., and Kraiselburd, E. (2006). Selective attention toward female secondary sexual color in male rhesus macaques. *American Journal of Primatology*, 68(7):738–744.

Wang, P. (2011). An LCD monitor with sufficiently precise timing for research in vision. *Frontiers in Human Neuroscience*, 5.

Wang, X. J. (2010). Neurophysiological and computational principles of cortical rhythms in cognition. *Physiological Reviews*, 90:1195–1268.

Wang, Y. and Dragoi, V. (2015). Rapid learning in visual cortical networks. *eLife*, 4.

Wang, Y., Iliescu, B. F., Ma, J., Josic, K., and Dragoi, V. (2011). Adaptive changes in neuronal synchronization in macaque V4. *Journal of Neuroscience*, 31(37):13204–13213.

Wässle, H. and Boycott, B. B. (1991). Functional architecture of the mammalian retina. *Physiological Reviews*, 71(2):447–80.

Wässle, H., Grünert, U., Röhrenbeck, J., and Boycott, B. B. (1989). Cortical magnification factor and the ganglion cell density of the primate retina. *Nature*, 341(6243):643–646.

Werner, A., Sharpe, L. T., and Zrenner, E. (2000). Asymmetries in the time-course of chromatic adaptation and the significance of contrast. *Vision Research*, 40(9):1101–1113.

Westerberg, J. A., Cox, M. A., Dougherty, K., and Maier, A. (2019). V1 microcircuit dynamics: altered signal propagation suggests intracortical origins for adaptation in response to visual repetition. *Journal of Neurophysiology*, 121(5):1938–1952.

Whittington, M. A., Cunningham, M. O., LeBeau, F. E., Racca, C., and Traub, R. D. (2011). Multiple origins of the cortical gamma rhythm. *Developmental Neurobiology*, 71:92–106.

Whittington, M. A., Traub, R. D., and Jefferys, J. G. (1995). Synchronized oscillations in interneuron networks driven by metabotropic glutamate receptor activation. *Nature*, 373:612–615.

Wiggs, C. L. and Martin, A. (1998). Properties and mechanisms of perceptual priming. *Current Opinion in Neurobiology*, 8(2):227–233.

Wissig, S. C. and Kohn, A. (2012). The influence of surround suppression on adaptation effects in primary visual cortex. *Journal of Neurophysiology*, 107(12):3370–3384.

Woloszyn, L. and Sheinberg, D. L. (2012). Effects of long-term visual experience on responses of distinct classes of single Units in inferior temporal cortex. *Neuron*, 74(1):193–205.

Womelsdorf, T., Lima, B., Vinck, M., Oostenveld, R., Singer, W., Neuenschwander, S., and Fries, P. (2012). Orientation selectivity and noise correlation in awake monkey area V1 are modulated by the gamma cycle. *Proceedings of the National Academy of Sciences*, 109:4302–4307.

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

Xing, D., Yeh, C.-I., Burns, S., and Shapley, R. M. (2012). Laminar analysis of visually evoked activity in the primary visual cortex. *Proceedings of the National Academy of Sciences*, 109(34):13871–13876.

Xing, D., Yeh, C.-I., and Shapley, R. M. (2010). Generation of black-dominant responses in V1 cortex. *Journal of Neuroscience*, 30:13504–13512.

Yeh, C. I., Xing, D., and Shapley, R. M. (2009). Black responses dominate macaque primary visual cortex V1. *Journal of Neuroscience*, 29(38):11753–11760.

Yoshida, M. and Isa, T. (2011). Express saccade without V1. *Neuroscience Research*, 71:e144.

Yu, J. and Ferster, D. (2013). Functional coupling from simple to complex cells in the visually driven cortical circuit. *Journal of Neuroscience*, 33(48):18855–18866.

Yu, Y., Schmid, A. M., and Victor, J. D. (2015). Visual processing of informative multipoint correlations arises primarily in V2. *eLife*, 4.

Zaidi, Q., Ennis, R., Cao, D., and Lee, B. (2012). Neural locus of color afterimages. *Current Biology*, 22(3):220–224.

Zeitler, M., Fries, P., and Gielen, S. (2006). Assessing neuronal coherence with single-unit, multi-unit, and local field potentials. *Neural Computation*, 18(9):2256–2281.

Zeki, S. M. (1978). Uniformity and diversity of structure and function in rhesus monkey prestriate visual cortex. *Journal of Physiology*, 277(1):273–290.

Zhaoping, L. (2019). A new framework for understanding vision from the perspective of the primary visual cortex. *Current Opinion in Neurobiology*, 58:1–10.

Zhu, M. and Rozell, C. J. (2013). Visual nonclassical receptive field effects emerge from sparse coding in a dynamical system. *PLoS Computational Biology*, 9(8):1–15.

Zurawel, G., Ayzenshtat, I., Zweig, S., Shapley, R., and Slovin, H. (2014). A contrast and surface code explains complex responses to black and white stimuli in V1. *Journal of Neuroscience*, 34(43):14388–14402.

Zweig, S., Zurawel, G., Shapley, R., and Slovin, H. (2015). Representation of color surfaces in V1: edge enhancement and unfilled holes. *Journal of Neuroscience*, 35(35):12103–12115.

List of Figures

List of Tables

Acknowledgements

Throughout the projects described in this dissertation, I have received a great deal of support. I want to thank my supervisor, Prof. Pascal Fries, for inspiration, methodological rigor, encouragement and freedom to pursue my own ideas. I am grateful to my promotor, Prof. Kössl, for his good-natured support and for agreeing to co-supervise a student with a research topic rather different from his own. I want to thank the two other members of my PhD committee, Prof. Wolf Singer and Dr. Michael Schmid, for continuous mentorship far exceeding the official job description, and providing opportunities that literally made my projects possible. For similar reasons, I also greatly appreciate supervison by and collaboration with Dr. Martin Vinck. Prof. Peter de Weerd and Prof. Onur Güntürkün sparked my interest in neuroscience. I thank all my mentors for sharing their complementary views on the mysterious workings of the brain and the mind.

All of my mentors came embedded in inspiring communities. In the context of this thesis, I am particularly grateful for the collaboration with my co-authors Cem Uran, Mark Roberts and Eric Lowet, as well as the Singer, Schmid and Fries lab. Inspired by Joscha Schmiedt, I will attempt a roughly historical order. Thank you Iris Grothe, Giorgos Spyroupoulos, Marieke Schölvinck, Jarrod Dowdall, Katharine Shapcott, Kleopatra Kouroupaki, Joscha Schmiedt, Liane Klein, Gareth Bland, Sylvia van Stijn, Will Barnes, Yiling Yang, Elena Psarou, and Benjamin Stauch, for collaborations past and present, and generous support. I want to stress invaluable support by the technicians of all labs: Rasmus Roese, Hanka Klon-Lipok, Sabrina Wallrath, Julia Hoffmann, and Nanni (Marianne) Hartmann. Thanks to our animal caretakers, in particular Andre Marron and Clemens Sommer, and vetenarians Christiane Kiefert, Christa Tandi and Alf Theisen, for excellent animal care and support. The IMPRS for Neural Circuits provided a welcoming community, great learning opportunities and financial support. I also want to thank Dr. Charlotte Koenen for proof-reading. All remaining errors, typographical or otherwise, are entirely my own.

Personal (Unprofessional) Acknowledgements

A remarkable number of people supported me in ways that are not always easily described. This list is largely unordered, very likely incomplete, and certainly not doing full justice to the people on it. I want to thank

- Drs Chris Lewis, Gareth Bland, Gustavo Rohenkohl and the rest of the typical yet ever-changing evening crowd for general discussions of all things science and the Binding moments at the Trinkhalle.
- To quote Ayelet Landau: "The women of the ESI."
- Giorgos Spyropoulus for a remarkably extensive vocabulary and an equally remarkable perspective on life.
- Dr Nina Merkel and Sylvia van Stijn, for countless jogging rounds that helped keep me sane (probably). Also for emotional support cat pictures.
- Clemens Sommer and Andre Marron for support, a sense of humor and the occasional hug as required.
- Hanka Klon-Lipok for her expert advice, and her love and dedication to her profession that is beyond description.
- Cem Uran for the best T-shirts, Sissi Tzanou for the best chocolate cookies, and Will Barnes for the best salsa.
- My office mates Jianguang Ni, Yufeng Zhang, and Ben Stauch, for a good atmosphere and enlightening conversation.
- Drs Andreea Lazar and Eric Lowet, for honesty, sharp insight, and friendship.
- Dr Pilar Cossio, for her bright mind, warm heart and a sea urchin shell that was much bigger and deeper in the water than it appeared.
- Liane Klein for photography that creates lasting memories, great food, patience, and her wisdom that if I am too sick to walk I shouldn't be working.
- Renata Vajda for books worth reading, her courage, her heart, and being who she is.
- Dr Katharine Shapcott, for her theories, her enthusiasm, and lasting friendship.
- Drs Charlotte Koenen and Clara Quetscher, for understanding that only common suffering can create, good cheer and RuPaul's Drag Race marathons.
- My entire family, for their trust and unwavering support.

I couldn't have done it without you.

Declaration

I declare that I have written this thesis independently and without any use of resources other than the ones quoted in the thesis. The thesis has not been used as an examination record in another course of studies. This is a true copy of the thesis, including any required final revisions, as accepted by my supervisors. I understand that my thesis may be made electronically available to the public.

Alina Peter
Frankfurt am Main
October 2019