RECURRENT NETWORK DYNAMICS MODULATE ORIENTATION TUNING IN PRIMARY VISUAL CORTEX

by

MARIA DEL MAR QUIROGA

A Dissertation submitted to the

Graduate School-Newark

Rutgers, The State University of New Jersey

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Behavioral and Neural Sciences

written under the direction of

Bart Krekelberg

and approved by

Dr. Jorge Golowasch		
U -		

Dr. Denis Pare

Dr. Farzan Nadim

Dr. Adam Kohn

Dr. Bart Krekelberg

Newark, New Jersey

May, 2015

©2015

Maria del Mar Quiroga

ALL RIGHTS RESERVED

ABSTRACT OF THE THESIS

RECURRENT NETWORK DYNAMICS MODULATE ORIENTATION TUNING IN PRIMARY VISUAL CORTEX

By MARIA DEL MAR QUIROGA

Thesis Director:

Bart Krekelberg

Adaptation refers to the phenomenon that sensory neurons are affected not only by their immediate input, but also by the sequence of preceding inputs. In visual cortex, for example, neurons change their preferred orientation after exposure to an oriented stimulus. This adaptation is traditionally attributed to biophysical changes in the intrinsic properties of neurons and their connections to other neurons (i.e. 'plasticity'). In this thesis, however, we propose that some effects of adaptation on neural responses instead reflect natural consequences of response dynamics in recurrent neural networks.

In Chapter 2 we used computational modeling to show that recurrently connected neurons can be surprisingly slow to respond to changing environments, even when they have fast intrinsic dynamics. Consequently, adaptation effects on a time-scale of hundreds of milliseconds arose naturally in a recurrent model of sensory processing without plasticity. We showed that these adaptation effects match previously reported changes in orientation tuning in cat and monkey primary visual cortex.

In addition, the model predicted a novel short-term change in orientation perception. In Chapter 3 we used quantitative psychophysics to test this prediction in human observers. These behavioral data provided additional support for the model.

In Chapter 4 we explored the dynamics of neural responses to rapidly presented sequences of oriented gratings, using electrophysiological techniques to record the activity of neurons in primary visual cortex. We directly compared the recorded neural responses to the responses of the model presented in Chapter 2. We showed that some, but not all, of the changes in the response due to adaptation can be explained by recurrent connectivity within the network. We also compared the neural responses to a linear summation model proposed by Benucci et al. (2009). We found that the linear prediction failed to explain the dynamic responses, suggesting that non-linear, recurrent interactions make important contributions to the representation of orientation.

Through the combination of computational, behavioral, and electrophysiological techniques, our work highlights the complexity of neural tuning properties that recurrent networks generate in dynamic sensory environments. Further, our work offers a path forward to measure and understand the complexity of recurrent neural networks in the alert brain.

Acknowledgements

I express my deepest gratitude and appreciation to my thesis advisor, Bart Krekelberg, for being the best mentor I could have ever hoped for. All of my future bosses will pale in comparison. He has served a multifaceted role as my advisor, colleague, pro-bono therapist, counselor, personality builder, family, landlord, and best friend (only after Adam and Olga left, of course). He has played a fundamental role in helping me become not only a better scientist, but also a better person.

I thank my thesis committee for their time and effort spent in improving my work. I specifically thank Jorge and Farzan for teaching the most interesting and enjoyable courses I've taken; and Denis for his personal support in addition to academic, including the best academic compliments I have ever received. In both these regards, a special thanks goes out to Jim Tepper as well. A heartfelt thanks to Ian Creese for getting me refugee status at CMBN.

I thank my fellow lab members and colleagues for their continued support and complicity. Thanks to Jess for giving me her shoulder to cry on whenever I needed it. Thanks to Anne for teaching me so much and so well, and both her and RAF for taking such good care of my beloved furry friends.

I thank my dearest friend Olga for always being there to fight me, yell at me, hug me, swim with me, dance with me, or braid my hair; whatever she thought I needed at the moment. I thank my life partner, Adam Morris, and my family for their unconditional support, sympathy, company, love and affection. I dedicate this thesis to them.

Table of Contents

ABSTRAC	CT OF THE THESIS	ii
Acknowled	lgements	iv
Table of Co	ontents	vi
List of Illus	strations	ix
Chapter 1.	Introduction	1
1.1 T	he representation of orientation in primary visual cortex	4
Chapter 2.	Adaptation without plasticity	8
2.1 In	troduction	8
2.2 M	lethods	9
2.2.1	From population responses to tuning curves	13
2.2.2	Parameter estimation and plausibility scores	14
2.2.3	Experimental data	17
2.3 R	esults	19
2.3.1	Responses to a single visual stimulus	21
2.3.2	Comparison to experimental data	28
2.4 D	iscussion	32
Chapter 3.	Neural population dynamics change perceived orientation	34
3.1 In	troduction	34
3.2 M	lethods	35

3.2.1	Participants	. 35
3.2.2	Apparatus	. 36
3.2.3	Visual Stimuli	. 36
3.2.4	Procedure	. 38
3.2.5	Trial presentation	. 40
3.2.6	Data analysis	. 41
3.3 Res	sults	. 42
3.3.1	Experiment 1	. 42
3.3.2	Experiment 2	. 45
3.4 Dis	scussion	. 48
Chapter 4.	Dynamic response properties in primate V1	. 51
4.1 Int	roduction	. 51
4.2 Me	thods	. 52
4.2.1	Electrophysiology	. 52
4.2.2	Visual stimuli	. 53
4.2.3	Procedure	. 54
4.2.4	Data analysis	. 54
4.3 Res	sults	. 57
4.4 Dis	scussion	. 71

Chapter	5.	Reverse connectivity simulations	76
5.1	Intro	oduction	76
5.2	Met	thods	77
5.3	Res	ults	79
5.3.	1	Motivation	79
5.3.	.2	Proof of principle	80
5.3.	.3	Robustness to noise	84
5.4	Dise	cussion	84
Chapter	6.	General discussion	87
6.1	Con	nclusions	93
Referen	ces		95

List of Illustrations

Figure 1 - Orientation selective V1 neuron, modified from Hubel and Wiesel (1962) 5
Figure 2 - Heuristic plausibility scoring functions, estimated from the experimental
literature15
Figure 3 - Response dynamics in a recurrent model of orientation selectivity
Figure 4 - Effect of parameters or context on the dynamics. Left panels correspond to the
C-model, right panels to the M-model
Figure 5 - Population responses to a horizontal grating after it was preceded by an adapter
orientation
Figure 6 - Recurrent network dynamics generate repulsive shifts in tuning in an
adaptation protocol
Figure 7 - The recurrent network models capture the magnitude and time course of tuning
curve shifts found in cat (Felsen et al., 2002) and macaque (Patterson et al., 2013)
V1
Figure 8 - Tuning curve shifts and intracortical connection profiles for the C-model (a),
the M-model (b), and a third model with slow dynamics (c)
Figure 9 - Size and duration of the shifts in preferred orientation for the slow model
shown in Figure 8 c
Figure 10 - Psychophysical measurement of short-term adaptation effects in humans 43

Figure 11 - Difference in PSEs between 20 and -20° adapting conditions, as a function of
test duration
Figure 12 - Task and group data for Experiment 2 46
Figure 13 - Mean shift across subjects as a function of the duration of the test stimulus. 48
Figure 14 - Extract of an example trial from experiment
Figure 15 – Shifts in the MUAe tuning curves for 30 array electrodes recorded in one session
Figure 16 - Principal components analysis for the multi-unit activity envelope (MUAe) of 30 simultaneously recorded electrodes with significantly tuned tuning curves 61
Figure 17 - Principal components analysis for 18 C-model units with equally spaced preferred orientations
Figure 18 - Obtaining the event related average response of one electrode
Figure 19 - Event related average response to each orientation relative to preferred for 30 well-tuned simultaneously recorded electrodes
Figure 20 - Event related responses to jumps in orientation, aligned to the onset of the second grating in the pair (at 0 ms, response seen after ~60 ms latency)
Figure 21 - Summing two responses to a single orientation does not explain the response to jumps in orientation – one recording session
Figure 22 - Summing two responses to a single orientation does not explain the response to jumps in orientation – average over 7 recording sessions

Figure 23 - Persistence in event related responses when followed by a blank screen 71
Figure 24 - Examples of different patterns of shifts for simultaneously recorded single
units in V1
Figure 25 - The model that best matches the response also results in a close match to the
connectivity profile, but is not necessarily the closest model to the "real" model in
parameter space
Figure 26 - Models that match the response of the experimental data don't necessarily
have similar connectivity profiles
Figure 27 - Adding noise to the C-model does not impair the ability to find the model
with the most similar connectivity profile

Chapter 1. Introduction

Recurrent networks are ubiquitous throughout the brain. Even so, much of the research in sensory processing being done today focuses on feed-forward processing of information and ignores the contribution of locally recurrent feedback. One reason for focusing mostly on feed-forward signals is that it is often hard to have an intuitive understanding of behaviors that can emerge from nonlinear, recurrently connected networks, so it is difficult to come up with testable hypotheses. Despite the relative success of feed-forward networks in explaining many behaviors observed in highly controlled experimental settings, there are several examples of behaviors that cannot be accounted for, and that can on the other hand be well explained by means of recurrent connections. One such example is that of persistent activity in prefrontal cortex, thought to be related to working memory processes. Compte et al. (2000) were able to account for experimentally measured persistent activity in dorsolateral prefrontal cortex with a recurrent network model similar to the one we present in Chapter 2.

In sensory regions, persistent activity would naturally lead to interactions in the coding of successively presented sensory stimuli. Feed-forward processes would update quickly to signal the current stimulus, while reverberation within the recurrent connections may still hold information about previous stimuli. In this work we focused particularly on the effect of recent stimulus history on neural representations, often called 'adaptation'. There are many examples in various areas of neuroscience that suggest an important role

of recurrent connections when responding to stimuli that are flanked in space or time by other stimuli. For example, recent stimulus history has been shown to produce nonlinear behavior in stereoscopic depth selectivity by Duong et al. (2011), which could not be explained by the current standing feed-forward model. In the auditory domain, Asari and Zador (2009) found that recent auditory context can affect the membrane potential of neurons in primary auditory cortex for up to four seconds after the end of the auditory stimulus. Ulanovsky et al. (2004) further suggested the existence of multiple timescales for encoding auditory stimulus history, which could be explained by different involvements of the recurrent connections for different types of auditory stimuli. Thus, further consideration of recurrent connectivity in the brain could explain many neural phenomena that are not readily accounted for by the purely feed-forward information processing in the brain.

Adaptation in early visual areas, sometimes associated with visual aftereffects such as the tilt aftereffect (TAE) (Gibson and Radner, 1937) or the motion aftereffect (MAE) (Addams, 1834), has commonly been attributed to plasticity mechanisms (e.g. changes in the effective connectivity of the recurrent connections) (for a review, see Kohn, 2007). Similar aftereffects, with similar proposed neural mechanisms have been found in higher visual areas with face gender aftereffects (FAE) (Zhao et al., 2011). This 'pairing' of adaptation and plasticity have led many researchers to use both terms interchangeably (for an example, see Dragoi et al., 2001). To clarify, in this work we use the term 'adaptation' in a mechanistically-neutral way to describe a change in either the observer's perception (i.e. after-effects) or neural behavior. We use the term 'plasticity', on the other

hand, to refer to a specific biophysical mechanism that involves a change in the intrinsic properties of a neuron, or a change in the strength of the synaptic connection to other neurons. Such changes have the *potential* to cause changes in the behavior of neurons and/or perception. In this sense some, but not all adaptation, is due to plasticity. Indeed, as we will show, not all instances of altered behavior are due to changes in the mechanisms of the system. As a dynamical system, the brain changes over time as its default state.

Our work investigates the effects of recent stimulus history and non-linear network dynamics in the particular domain of visual orientation selectivity, but findings could extend by analogy to any recurrently connected network in the brain. Our primary hypothesis is that many of the effects of adaptation previously attributed to mechanisms of plasticity are instead (or perhaps in addition) attributable to the influence of recurrent dynamics on the tuning properties of single neurons.

The work described in this thesis includes three distinct approaches to this problem: computational modeling, psychophysics, and multi-electrode recordings of neural activity in primary visual cortex of macaques. We start with a theoretical analysis in which we show that a recurrent model of orientation selectivity is able to reproduce many of the experimentally observed adaptation effects without the need for plasticity (Chapter 2). We further validated this model using a behavioral study in which we tested psychophysically some of its predictions on perception (Chapter 3), and by measuring neural responses in primary visual cortex (Chapter 4). We use the unique insight gained by the combination of the three methodologies to advance towards a proof of principle for a novel way to use the model as a tool to infer networks properties that are difficult to measure directly (Chapter 5).

We will start with a broad overview of orientation selectivity in primary visual cortex, and we will introduce the more specific background to each part of this project in the Introduction sections of each chapter.

1.1 The representation of orientation in primary visual cortex

Neurons in primary visual cortex (V1) are the first in the visual pathway to show orientation selectivity in their firing patterns. In their pioneering work, Hubel and Wiesel (1962) proposed that selective pooling of inputs from the lateral geniculate nucleus (LGN) accounts for this selectivity. They found that LGN neurons have circular on/off center-surround receptive fields. This means that some LGN neurons respond to a circle of light surrounded by a dark annulus (on-center cell), and others to a circle of dark surrounded by a light annulus (off-center cell). If several on-center cells are aligned on an axis, and they are all synaptically connected to a V1 neuron, then this V1 neuron becomes orientation selective (Figure 1).

This feed-forward model explains an impressive number of experimental findings, but not all known characteristics of V1 tuning curves, such as their contrast invariance and effects of adaptation (Ferster and Miller, 2000). This is not surprising, given that more than half of the inputs to V1 neurons arise from neighboring cortical cells (Chung and Ferster, 1998). More recently, it has become more common to consider recurrent networks and their capabilities for information processing.



Figure 1 - Orientation selective V1 neuron, modified from Hubel and Wiesel (1962). Four on-center LGN neurons converge onto a V1 neuron, which will have a preferred orientation of around 30°, consistent with the axis on which the on-center cells are aligned.

One paradigmatic model of orientation selectivity that included recurrent interactions was proposed by Somers et al. (1995). It was carefully parameterized based on biophysical measures of neuronal properties from the literature and was able to account for many experimentally observed properties of orientations selectivity, such as sharpening of tuning curves and contrast invariance of responses. Even more impressive was its ability to predict the effect of some pharmacological manipulations, such as Sillito et al.'s (1980) blocking of inhibition in cortex by cooling, and even blockage of inhibition in a single cell (Nelson et al., 1994). The model of Somers et al. (1995) provided a useful avenue for studying the influence of recurrent connectivity on orientation selectivity, but it had a very high number of free parameters and so was computationally expensive to run simulations and to fit it to data. Fortunately, Carandini and Ringach (1997) noticed that the model could be reduced to one single differential equation for each model neuron, while still quantitatively replicating the explanatory power of the more detailed model. They did, however, find that the model did not respond appropriately to plaid stimuli (i.e. grating stimuli that contain two or more superimposed orientations), given its tendency to converge to one peak of activity if the two orientations were close (~30° apart), and to push the two peaks away from each other if the two orientations were further from each other (~60° apart).

Teich and Qian (2003) subsequently refined the model of Carandini and Ringach (1997) and in doing so solved some of its problems, like the rising of spurious activity. Teich and Qian (2003) were also the first to show that this model could explain effects of learning and adaptation on orientation selective neurons by modifying the strength of the intracortical connections in an orientation specific manner, simulating plasticity. As we will explain in Chapter 2, however, for short-term adaptation their model would have generated the same effects even before they modified the intracortical interactions.

The model of orientation selectivity presented in this thesis builds primarily on the models of Somers et al., (1995), Carandini and Ringach (1997), and Teich and Qian (2003). Our primary contribution is to investigate the dynamics of the model rather than the properties of the steady state responses. As we show in the chapters, these dynamics are complex, and can account for a number of findings in behavioral studies as well as

the neural response properties of primary visual cortex. As we will argue in Chapter 6, these dynamics are particularly relevant under natural viewing circumstances, when a recurrently connected network rarely reaches a steady state. By building on these models, and combining computational, behavioral, and electrophysiological techniques, our work has the potential to provide much needed insight into the complex neural dynamics generated by recurrent connectivity.

Chapter 2. Adaptation without plasticity

2.1 Introduction

Neurons in primary visual cortex (V1) are strongly tuned to the orientation of visual stimuli within their receptive fields (Hubel and Wiesel, 1962). Some debate persists surrounding how neurons acquire this tuning (Ferster and Miller, 2000), including the relative influences of feed-forward versus recurrent inputs (Sompolinsky and Shapley, 1997; Priebe and Ferster, 2008). However, it is widely accepted that initial tuning is determined by selective pooling of inputs from the lateral geniculate nucleus (LGN) and perhaps further modified (sharpened, made contrast invariant) by the recurrent cortico-cortical connections that account for up to half of the inputs to V1 neurons (Chung and Ferster, 1998).

Recent studies show that orientation preference is altered by exposure to oriented stimuli (Muller et al., 1999; Dragoi et al., 2000; Dragoi et al., 2002; Felsen et al., 2002; Kohn, 2007; Wissig and Kohn, 2012; Patterson et al., 2013). These alterations in neural responses have been linked to perceptual phenomena such as the tilt after-effect (TAE), where adaptation produces a 'repulsive' shift in orientation perception (i.e. perception of orientation is biased away from that of the adapter stimulus) (Gibson and Radner, 1937). Remarkably, neural adaptation can be observed even when the adapting stimulus is presented for only a fraction of a second (Muller et al., 1999; Felsen et al., 2002), suggesting a very rapid form of reorganization. Analogous adaptation effects are found in all sensory brain areas and are thought to provide a functional benefit by enhancing discriminability of stimuli that are prevalent in the environment (Muller et al., 1999; Krekelberg et al., 2006a; Kristjansson, 2011), or by increasing detectability of rare stimuli (Clifford et al., 2001; Dragoi et al., 2002). Crucially, adaptation and its associated benefits are commonly thought to arise from changes in the intrinsic properties of neurons or their inter-connections, such as the efficacy of synaptic inputs (i.e. 'plasticity') (Felsen et al., 2002; Teich and Qian, 2003). Here, however, we evaluated and confirmed a more parsimonious explanation: that short-term adaptation (and therefore its potential associated benefits) are emergent properties of processing within a recurrent neural network.

2.2 Methods

We implemented a recurrent model of orientation selectivity following the work of Somers et al. (1995); Carandini and Ringach (1997) and Teich and Qian (2003). The model consisted of coupled differential equations describing the membrane potential of N orientation selective primary visual cortex neurons (Figure 3). As a whole the network represents a cortical hypercolumn (i.e. all N spatial receptive fields overlap) and all N neurons are identical except for their preferred orientation. The preferred orientations of the N model neurons were evenly distributed across all orientations (-90° to 90°), and we labelled each model neuron based on its preferred orientation θ . Each neuron was modelled as a single passive voltage compartment, whose membrane potential over time $V^{\theta}(t)$ obeyed the differential equation:

$$\tau \frac{dV^{\theta}}{dt} + V^{\theta} = V^{\theta}_{\lg n} + V^{\theta}_{cortex} \tag{1}$$

where τ is the membrane time constant, V_{lgn}^{θ} the synaptic potential generated by the thalamocortical inputs to the model neuron and V_{cortex}^{θ} is the net synaptic input to the neuron from its cortical neighbors (Carandini and Ringach, 1997; Teich and Qian, 2003).

We obtained an instantaneous firing rate for the neuron with preferred orientation θ at time *t* according to the equation:

$$R^{\theta}(t) = \alpha \max\left(V^{\theta}(t), 0\right)$$
(2)

where $V^{\theta}(t)$ is the voltage of the neuron with preferred orientation θ at time t (see Figure 3c). This results in a mean firing rate model, with α the corresponding gain factor (i.e. increase in firing rate [Hz] for a 1 mV increase in the membrane potential). For simplicity, the neurons have zero spontaneous firing rate and the membrane potential was measured relative to the spike threshold. In other words, a positive membrane potential results in spikes, whereas a negative membrane potential reflects a modulation for the model neuron below the spiking threshold, hence the zero on the right term of equation (2).

Our implementation is similar to the work of Carandini and Ringach (1997) and Teich and Qian (2003) except that we used von Mises functions to represent the thalamocortical input as well as the lateral connectivity in the network. Using this differentiable and circular function instead of the non-circular (truncated) Gaussians of previous studies, improved the numerical behavior of the network and removed artefacts from the dynamics. For each model cortical neuron with preferred orientation θ , the input from LGN was a function of stimulus orientation ω and contrast c:

$$V_{\lg n}^{\theta}(\omega,c) = cJ_{\lg n}f(\omega \mid \theta,\kappa_{\lg n}), \qquad (3)$$

Where $J_{\lg n}$ represents the strength of the input and $f(\omega | \theta, \kappa_{\lg n})$ is the von Mises distribution with period π , mean θ and concentration $\kappa_{\lg n}$ (inversely related to the width - Figure 3a):

$$f(x \mid \mu, \kappa) = \frac{e^{\kappa \cos(2(x-\mu))}}{2\pi I_0(\kappa)}$$
(4)

 $I_o(\kappa)$ is the modified Bessel function of order zero. As in previous implementations (Carandini and Ringach, 1997; Teich and Qian, 2003) we did not incorporate separate units to represent excitatory and inhibitory cells, but rather gave each model neuron the ability to produce both net excitation and inhibition. Experimental studies have found that the probability of connection between two orientation selective neurons is not uniform. Instead, the probability of connection is highest for neurons with similar preferred orientation and decreases with increasing difference in preferred orientation (Michalski et al., 1983; Ko et al., 2011). Based on these findings, we modelled both the excitatory (E^{θ})

and inhibitory (I^{θ}) connection profiles to the model neuron with preferred orientation θ as von Mises distributions in orientation space (see equation (4)):

$$E^{\theta}(\varphi) = f(\varphi \mid \theta, \kappa_{E})$$
⁽⁵⁾

$$I^{\theta}(\varphi) = f(\varphi \mid \theta, \kappa_{I})$$
⁽⁶⁾

When the inhibitory connection profile was (marginally) broader than the excitatory one ($\kappa_E > \kappa_I$), the connection profile ($F^{\theta}(\varphi)$, Figure 1b) had the typical 'Mexican hat' shape:

$$F^{\theta}(\varphi) = J_{cortex} \left(E^{\theta}(\varphi) - r_{IE} I^{\theta}(\varphi) \right)$$
(7)

In this recurrent connection profile J_{cortex} represents the strength of the cortical connections, and r_{IE} the ratio of the strength of inhibition to the strength of excitation.

At each point in time, the membrane potential for the neuron with preferred orientation θ induced by the recurrent input is a weighted sum of the firing rates of all neurons in the network (equation (2)):

$$V_{cortex}^{\theta}(t) = \sum_{\varphi} F^{\theta}(\varphi) R^{\varphi}(t)$$
(8)

The model and all simulations were implemented in MATLAB version 8.2.0.701 and we solved the model numerically using *ode45*, an adaptive time step Runge-Kutta method.

The number of model units in the network (N) was set to 256 and the stimulus contrast c was set to 50%. Source code is available on request.

Given our main purpose of studying temporal dynamics induced by recurrent connectivity, and unlike previous work (Carandini and Ringach, 1997; Teich and Qian, 2003), we explored model responses at every time point during the simulation as opposed to only the final steady state responses. This allowed a direct comparison of the simulations with experimental data.

2.2.1 From population responses to tuning curves

The output of the model is the response of all neurons to a single stimulus (equations (1) and (2)). Lining up all model neurons on the x-axis according to their preferred orientations and plotting the response of each one on the y-axis, we obtain what we call the population response (Figure 3d). Electrophysiology experiments, on the other hand, typically report tuning curves of individual neurons: the response of a single neuron to all possible stimulus orientations. To obtain tuning curves for our model neurons, we simulated the population response to a range of stimulus orientations, and extracted from each of these the response of one specific model neuron.

The tuning curves presented in Figure 6b were obtained following either a standard or an adaptation tuning protocol (Figure 6a). Mimicking the standard procedure in electrophysiological experiments, for the standard protocol we presented one test orientation per trial and averaged the firing rate over the duration of the test stimulus. To preclude systematic interactions beyond the measurements in one trial, the response of

each neuron was reset to the spontaneous firing rate at the end of the trial (experimentally this would correspond to a long inter-trial interval). In the adaptation protocol, on the other hand, we preceded each test orientation with a fixed adapting orientation (-20° for 20 ms; Figure 6a).

2.2.2 Parameter estimation and plausibility scores

The three models considered in this paper (C-model, M-model, Slow-model) all exhibit response properties (e.g. tuning width, time-to-peak) that resemble those observed experimentally in V1. We first performed an exhaustive search to find free parameters that generated single neuron responses we considered to be plausible based on the literature. For instance, for some model parameters neurons had peak firing rates of 200 spikes per second. Such responses are unusual in experimental recordings, hence we assigned low plausibility scores to models that predicted such outcomes. An analogous scoring heuristic was applied based on single neuron's time to peak response, and single neuron tuning widths (Figure 2). To obtain the score functions presented in Figure 2 we fitted kernel probability distributions to histograms of experimental data, using speciesspecific estimates where possible (DeValois et al., 1982; Gardner et al., 1999; Carandini and Ferster, 2000; Schummers et al., 2007; Wissig and Kohn, 2012). When more than one dataset was available for the same experimental measure, we first normalized and summed the histograms before fitting the kernel distribution. The resulting distributions were then scaled to span the interval [0 1]. To score each model we simulated the response to single gratings of all orientations, measured the time to peak, the peak response, and the tuning width, extracted the corresponding three scores from the three

heuristic score functions (Figure 2), and averaged them to get a final plausibility score between zero and one. Note that plausibility as used here only refers to single stimulus responses, measured outside the context of an adaptation protocol.



Figure 2 - Heuristic plausibility scoring functions, estimated from the experimental literature. **a**: Scoring function for tuning curve bandwidth based on anaesthetized cat recordings (Gardner et al., 1999; Carandini and Ferster, 2000). **b**: Tuning bandwidth score for anaesthetized monkey recordings (DeValois et al., 1982). **c**: Time-to-peak score based on anaesthetized cat recordings (Schummers et al., 2007). **d**: Peak firing rate score based on anaesthetized monkey recordings (Wissig and Kohn, 2012).

This exhaustive search for plausible models was limited to parameters with the following values (minimum:stepsize:maximum): $\alpha = [4:1:13]$ Hz/mV; $J_{lgn} = [7:1:16]$ mV/Hz; $\kappa_{Ign} = [0.5:0.2:2.3]; J_{cortex} = [0.9:0.1:1.8]$ mV/Hz; $r_{IE} = [1:0.02:1.18]; \kappa_{E} = [1:0.0$

[1.5:0.1:2.4]; $\kappa_I = [1:0.1:1.9]$; τ was fixed to 15 ms during this exhaustive search. For every one of the 10⁷ factorial combinations of these free parameters, we scored the resulting models according to the heuristic plausibility score and used the most plausible models (scores above 0.9) as initial estimates for an optimization procedure. The goal of this optimization was to find parameters that best fit a specific set of experimental data obtained in an adaptation paradigm.

For each plausible initial guess, we determined the eight free parameters that matched the experimental data set most closely in the least squares sense (*fmincon* in MATLAB). From this optimization, we selected the parameter sets with the lowest error and the highest plausibility score.

Optimization was done separately for data obtained in cat and macaque primary visual cortex. We refer to the best-fitting model for each of these data sets as the C-model and the M-model, respectively. The model fitting procedure was robust: we found large, contiguous parts of parameter space with a good match between the C-model and the data, as well as the M-model and the data. The time-to-peak, peak firing rate, and tuning width were well within the experimentally observed ranges for all model variants in the main text. The qualitative finding that recurrent networks generated repulsive shifts in tuning curves (Figure 6) was also robust: shifts in tuning curves were present, with different magnitudes and time courses, for all parameter sets that resulted in high plausibility scores.

For the C-model the eight free parameters were optimized to match adaptation in cat visual cortex (Felsen et al., 2002). This resulted in the following parameters: $\tau = 10.8$ ms; $\alpha = 10.6$ Hz/mV; $J_{lgn} = 9.57$ mV/Hz; $\kappa_{lgn} = 1.56$; $J_{cortex} = 1.71$ mV/Hz; $r_{IE} = 1.18$; κ_E = 1.59; $\kappa_I = 1.16$. For the M-model, which was optimized to match adaptation in anaesthetized macaque visual cortex (Patterson et al., 2013), the parameters were: $\tau = 8$ ms; $\alpha = 3.88$ Hz/mV; $J_{lgn} = 11.04$ mV/Hz; $\kappa_{lgn} = 0.47$; $J_{cortex} = 2.84$ mV/Hz; $r_{IE} = 1.24$; $\kappa_E = 1.12$; $\kappa_I = 0.56$.

For the model shown in Figure 8 we fixed τ to 15 ms and then performed a random search through the remaining seven dimensional parameter space. This search was not exhaustive; we stopped the multi-day search after finding several models with tuning curve shift time courses whose dynamics were on the order of seconds. The parameters of the exemplary slow model in Figure 8 were: $\tau = 15$ ms; $\alpha = 4$ Hz/mV; $J_{1gn} = 8$ mV/Hz; $\kappa_{Ign} = 0.5$; $J_{cortex} = 1.7$ mV/Hz; $r_{IE} = 1.14$; $\kappa_E = 2.2$; $\kappa_I = 1$, and its plausibility score was 0.4 if scored with cat plausibility, 0.6 if scored with anaesthetized monkey plausibility.

2.2.3 Experimental data

In the Results section we compare our model results to those of two datasets from the literature (Felsen et al., 2002; Patterson et al., 2013). We report their methods here briefly. Felsen et al. (2002) used single tungsten electrodes to record from neurons in primary visual cortex of anaesthetized cats. The grating stimuli were adjusted to match the optimal spatial frequency and receptive field size of the neuron under study. The

stimuli were pairs or sequences of gratings in which each grating was presented for 25 ms or 17ms. The time between two gratings was varied between 0 ms and 125 ms; the screen was blank during that period of time. We extracted the tuning curve shift data from their Figure 3B (based on the average tuning curve shift of 21 cells; reproduced here as Figure 7a).

We extracted the time courses from their Figure 4B, which has separate time courses for 25 ms and 17 ms grating duration for a total of 20 cells. We also redefined "interval" to be the time (in ms) that the screen was blank in between two oriented gratings. The model could be fit to either of these time courses (not shown) but for simplicity we determined a single average time course by averaging the two curves (shown here in Figure 7c) and used the average stimulus duration (20 ms) for model simulations. We use 'C-model' to refer to the best-fitting model for these data.

Patterson et al. (2013) used multi-electrode arrays to record from neurons in layers 2/3 and 4B of the parafoveal representation in primary visual cortex of the anaesthetized macaque. The grating stimuli were adjusted to cover the aggregate receptive field of the neurons recorded by the electrode array, their spatial frequency was fixed at 1 cycle/° and they drifted at 6.25Hz. Of main interest here were their experiments that used a brief adapter (400 ms) followed by a brief test (400 ms). Patterson et al. (2013) mapped the time course of tuning curve shifts by measuring the neural response in three separate epochs after test stimulus onset; early: 50-100 ms, mid: 100-200 ms, and late: 200-400 ms. We extracted the time course from their Figure 5C based on the average tuning curve shift of between 96 and 117 cells, and compare it to model results in our Figure 7d.

Because our model has no visual latency we adjusted the time epochs to be 20-70 ms (early), 70-170 ms (mid), and 170-370 ms (late), assuming a visual latency for the data in Patterson et al. (2013) of 30 ms. We use 'M-model' to refer to the best-fitting model for these data.

To simulate these experimental paradigms we presented pairs of adapt/test gratings to the model. To replicate the Felsen et al. (2002) data, the adapting and test orientations were varied factorially, choosing from 12 equally distributed orientations between -90° and $+90^{\circ}$. To replicate the Patterson et al. (2013) data, the adapter had one of 13 equally distributed orientations between 15° and 75° , and the test one of 20 equally distributed orientations between -90° and $+90^{\circ}$. From this data set we extracted the tuning curves for each adapter orientation. We also estimated tuning curves without an adapter. The tuning curve shift was then calculated as the difference between the orientation at the peak of the adapted and non-adapted tuning curves.

2.3 Results

We studied the dynamics of a recurrent network model of orientation tuning in primary visual cortex (Somers et al., 1995; Carandini and Ringach, 1997; Teich and Qian, 2003) (see Methods). The main properties and biological plausibility of this model have been documented in previous studies (Ben-Yishai et al., 1995; Somers et al., 1995; Carandini and Ringach, 1997). Briefly, the model consisted of a bank of V1 units, each with a different preferred stimulus orientation determined by their weakly-tuned feed-forward input from a thalamic layer (Figure 3a). Lateral connectivity for each neuron was modelled as the difference of two von Mises functions - one representing excitatory

connections and the other inhibitory; for appropriate parameter choices this results in a typical 'Mexican hat' profile (Figure 3b). The membrane potential of each model neuron was converted to an instantaneous firing rate using a piecewise linear function (Figure 3c).

The free parameters of the model were chosen to satisfy two criteria: (1) To produce plausible V1 response and tuning properties (peak firing rate, tuning bandwidth, and time-to-peak; see Methods and Figure 2), and; (2) To produce response dynamics that matched those of cat (Felsen et al., 2002) or macaque (Patterson et al., 2013) visual cortex observed experimentally (see Methods). We will refer to the cat and macaque parameterizations as the C-model and M-model, respectively. The results presented herein were derived from the C-model unless otherwise noted, but similar results were obtained with the M-model.



Figure 3 - Response dynamics in a recurrent model of orientation selectivity. **a**: Thalamocortical input to each model neuron for a horizontal grating (0° ; see inset) as a function of the neurons' preferred orientation. **b**: Intracortical connection profile for the neuron with preferred orientation 0° . **c**: Piecewise linear function that converts membrane potential to firing rate. **d**: Population response to a horizontal grating presented for 80 ms. Color represents time.

2.3.1 Responses to a single visual stimulus

When presented with a horizontal grating (0°) , the model initially showed a broadly distributed response across neurons with different preferred orientations (Carandini and

Ringach, 1997; Benucci et al., 2009). Over time, the response increased and sharpened, converging to a tuning width of \sim 32° (Figure 3d). This steady state response continued as long as the stimulus was present.



Figure 4 - Effect of parameters or context on the dynamics. Left panels correspond to the C-model, right panels to the M-model. **a:** Peak normalized firing rate after grating onset shown for networks with absent ($J_{cortex} = 0$; light grey curve), intermediate (C-model on left, M-model on right; medium grey curve), and strong (4 times the value used in the C-model and M-model respectively; black curve) recurrent connectivity. **b:** Response of the model neuron with preferred orientation 0° to the onset of a horizontal grating (time = 0) after it was preceded by one of three "adapter" orientations: 10° (purple), 60° (green), or 90° (red).

To investigate the role of recurrent connectivity on neural response dynamics, we modified the strength of the intracortical connections and examined the response to a grating. A purely feed-forward variant (recurrent inputs fixed at zero) reached steady

state after ~40 ms. Recurrent variants, in contrast, had convergence times that were substantially longer (150%) and increased with the strength of intracortical connectivity (Figure 4a, left panel). These effects were even stronger in the M-model (Figure 4a, right panel), and the model with four-fold stronger connectivity failed to converge entirely (black curve). Thus, response dynamics were not only determined by neurons' intrinsic properties (which were fixed in these simulations), but also by their connectivity.

Responses to sequential stimuli

Convergence time also depended on the network state induced by preceding sensory inputs. To illustrate this we show the response of neurons to sequential pairs of gratings (Figure 4b). The first grating (i.e. the 'adapter') was presented long enough that the network fully converged. Of key importance is what happened in the network when the stimulus was then switched to a new orientation (i.e. the 'test'). Interestingly, the network dynamics depended on the difference in orientation between the two successive gratings. For instance, when a 0° grating was preceded by a nearby orientation (10°), the neuron that preferred 0° gratings reached a steady state response 40 ms after the onset of the second grating (purple curve, Figure 4b, left panel). When preceded by a distant orientation (60° or 90°), however, the same neuron required ~50% more time to converge (green and red curves, Figure 4b, left panel). The M-model had even slower dynamics (Figure 4b, right panel).

The more nimble response for test stimuli near the adapter orientation is the somewhat trivial consequence of the fact that the adapter response closely resembles that of the test

and thus gives it a head start at the beginning of the test period. While this result is intuitive, we show below that these temporal interactions generate complex dynamics in the population and bias the representation of orientation when stimuli change rapidly.

We examined these dynamics by simulating the response to pairs of gratings (adapter/test) with different duration. After the onset of the test grating in a $-45^{\circ}/0^{\circ}$ pair with 80 ms presentation time per grating, the population activity declined from its peak at -45° and gradually shifted to re-establish itself around the neurons preferring 0° (Figure 5a). For pairs with similar orientations ($-20^{\circ}/0^{\circ}$), population activity resembled a travelling wave shifting between the presented orientations (Figure 5c). Finally, when the pairs were orthogonal, ($-90^{\circ}/0^{\circ}$; Figure 5d), activity collapsed around -90° as it resurfaced around 0°. Such dynamical patterns of network activity have been observed experimentally using voltage sensitive dye imaging in the cat visual cortex (Benucci et al., 2009).


Figure 5 - Population responses to a horizontal grating after it was preceded by an adapter orientation (color represents time after the onset of the horizontal stimulus). **a**: Response to an 80 ms presentation of a stimulus with orientation 0° , preceded by an 80 ms presentation of a stimulus oriented at -45° ($-45^{\circ}/0^{\circ}$ pair). The final population response to the second 0° stimulus (160 ms) peaks at the neuron with preferred orientation 0° . **b**: Response to a $-45^{\circ}/0^{\circ}$ pair with each grating presented for 20 ms. **c**: Response to a 20 ms $-20^{\circ}/0^{\circ}$ pair. **d**: Response to a 20 ms $-90^{\circ}/0^{\circ}$ pair. Short stimulus presentations bias the population response for some (**b**, **c**), but not other sequences (**d**).

Thus, for 80ms presentations, the population exhibited complex transitions from adapter to test but nevertheless represented orientation accurately by the end of the test period. Crucially, however, this was not the case for shorter presentation times because the transitions were interrupted midway. As a consequence, the peak of the population response at the end of the test stimulus presentation could be biased towards the neurons that preferred the adapter (Figure 5b). The size of this bias depended on the orientation difference between the adapter and the test grating in a non-monotonic fashion (compare panels b-d in Figure 5). This analysis shows that the response of a neuron in the recurrent network was not only determined by its immediate sensory input, but also the preceding inputs. In other words, the recurrent network demonstrated a property commonly termed adaptation.

To allow a direct comparison with experimental results (which typically report tuning curves of individual neurons rather than population responses), we simulated neural responses in a standard orientation tuning protocol and in an adaptation protocol (Figure 6a). In both cases we measured the tuning curve of a neuron whose nominal preferred orientation was 0° . As expected, the standard orientation tuning curve had a peak at the 0° orientation (Figure 6b, grey curve). In the adaptation paradigm, however, the head start in the response produced by the adapter (Figure 5) resulted in, somewhat counterintuitively, a shift of the tuning curve away from the adapter orientation (Figure 6b, black curve). Even though the two curves are peak-normalized, the shift in preferred orientation seems to come about as a result of an increased response in the opposite flank, rather than a reduced response in the adapted flank. This effect of adaptation on the far flank of neural tuning curves has been shown to be ubiquitous in primary visual cortex; particularly near pinwheel centers or on the border between orientation domains (Dragoi et al., 2001). This also suggests that our model's connectivity profile might resemble that

of a neuron in or near a pinwheel center; rather than a neuron within an iso-orientation domain.



Figure 6 - Recurrent network dynamics generate repulsive shifts in tuning in an adaptation protocol. **a**: Standard and adaptation orientation tuning protocols. **b**: Peaknormalized tuning curves for a neuron that prefers horizontal gratings obtained under the standard (grey) and adaptation (black) protocol. In the adaptation protocol, the peak of the tuning curve was at 3° . In other words, the -20° adapter (arrow) induced a repulsive shift in the tuning curve. **c**: Same as **b**, but not normalized.

In the adaptation literature a shift in orientation tuning such as the one shown in Figure 6b is known as a 'repulsive shift'. Such shifts have been found in a number of studies and

across different species (Muller et al., 1999; Dragoi et al., 2000; Dragoi et al., 2001; Dragoi et al., 2002; Felsen et al., 2002; Wissig and Kohn, 2012; Patterson et al., 2013) and are often interpreted as a sign of plasticity (Dragoi et al., 2000; Dragoi et al., 2001; Dragoi et al., 2002; Felsen et al., 2002; Teich and Qian, 2003). Our results however, raise the interesting alternative interpretation that these shifts reflect the natural dynamic behavior of a network of sensory neurons in a changing visual environment.

2.3.2 Comparison to experimental data

To explore this issue directly, we investigated whether the model could account quantitatively for the adaptation effects in two electrophysiological studies. Felsen et al. (2002) measured the size and time course of repulsive tuning curve shifts using sequences of very brief (17-25 ms) stimuli in cat V1 (see Methods; data re-plotted in Figure 7a and c). Remarkably, the C-model captured the direction, magnitude (Figure 7b), and time course (Figure 7c) of the shifts in that study. In a comparable study of macaque V1, Patterson et al. (2013) used longer adaptation and test stimuli (400 ms each) and found larger and longer-lasting (~200 ms) repulsive shifts than those of Felsen et al. (2002). Nevertheless, we found that recurrent network dynamics also accounted for these adaptation effects, albeit with somewhat different parameter values (i.e. the M-model; Figure 7d).



Figure 7 - The recurrent network models capture the magnitude and time course of tuning curve shifts found in cat (Felsen et al., 2002) and macaque (Patterson et al., 2013) V1. Grey background areas represent attractive shifts, and white background repulsive shifts. **a**: Tuning curve shifts averaged over all recorded neurons (preferred orientation set to 0°) in cat V1 (Felsen et al., 2002). **b**: Tuning curve shifts for the C-model neuron with preferred orientation 0°. **c**: Time course of tuning curve shifts for the cat data corresponding to (**a**) and the C-model. **d**: Time course of tuning curve shifts averaged over all recorded neurons (preferred orientation 0°. **c**: Time course of tuning curve shifts for the cat data corresponding to (**a**) and the C-model. **d**: Time course of tuning curve shifts averaged over all recorded neurons from macaque V1 (Patterson et al., 2013) and the M-model.

Interestingly, the main distinction between the C- and M-model parameters was that broader inhibitory lateral connections were required to reproduce the larger and longer duration shifts in the macaque data (Figure 8). While we cannot exclude the possibility of technical differences between the experiments (e.g. single electrode vs. array recordings, stimuli optimized for single neurons or for the average population), the requirement for broader inhibition could reflect true variation in connectivity between animal species, or recording sites. This suggests that an experimentally observed pattern and time course of tuning curve shifts could be used to infer a neuron's functional connectivity.



Figure 8 - Tuning curve shifts and intracortical connection profiles for the C-model (**a**), the M-model (**b**), and a third model with slow dynamics (**c**). The top panels show the magnitude of tuning curve shifts as a function of time after test onset. The tuning curves were obtained by averaging the response for the whole duration of the test stimulus. Tuning curve shifts on the order of several degrees and a time scale of seconds can be explained by a recurrent network model and do not require plastic network changes. The bottom panels show the intracortical connection profile for the neuron with preferred orientation 0°. The parameters of the slow model were chosen to generate slow but still V1-like response properties and tuning by a random search through all possible models, while keeping the neurons' membrane time constant $\tau = 15$ ms.

A wider parameter exploration (i.e. not tailored to fit any specific experimental data set) showed that the model could generate very large (~10°) and long-lasting (~1 s) tuning curve shifts (Figure 8c). Thus, these findings demonstrate that even if all neurons are fast when measured in isolation ($\tau = 15$ ms), a recurrent network can produce surprisingly slow and complex dynamics and adaptation effects without any contribution of plasticity.

Cortical strength (J_{cortex}) and ratio of inhibition to excitation (r_{IE}) are two model parameters that change the net excitation in the model. J_{cortex} does so by scaling the net excitation (as well as net inhibition) and r_{IE} by changing the balance between excitation and inhibition. Both these effects determine how slow the model is, and how large the shifts are, as observed in Figure 9.



Figure 9 - Size and duration of the shifts in preferred orientation for the slow model shown in **Figure 8 c**. Shown in red are the actual values of the slow model ($J_{cortex} = 1.7$ and $r_{IE} = 1.14$). **a:** Duration of the shifts (expressed as the last time point for which the shift is larger than 1°) as a function of J_{cortex} . **b:** Size of the shifts (expressed as the maximum shift attained) as a function of J_{cortex} . **c:** Duration of the shifts as a function of r_{IE} . **d:** Size of the shifts as a function of r_{IE} .

2.4 Discussion

We have shown that recurrent processing gives rise to repulsive shifts in tuning curves that mimic those observed experimentally in adaptation studies of primary visual cortex. Specifically, we showed that model V1 neurons shifted their preferred orientation away from that of an adapting stimulus for up to 1 second. This effect was a direct consequence of an interaction between the changing sensory environment and the dynamics of the processing network, and not from changes in the properties of individual neurons or their inter-connections. Thus, our results offer an alternative and parsimonious interpretation for many experimental phenomena previously attributed to rapid neural plasticity.

Nevertheless, intrinsic properties of single neurons undoubtedly do change during some adaptation protocols (Sanchez-Vives et al., 2000). This plasticity likely underlies response changes on a time scale of seconds to minutes (Kohn, 2007), which our model does not capture. By isolating the contribution of the recurrent connectivity, we provide a proof of principle: network dynamics generate adaptation-like phenomena on an ecologically relevant time scale of at least hundreds of milliseconds. More research is needed to investigate how these adaptation effects interact with those that arise from plasticity at different levels of the sensory hierarchy (Sclar et al., 1989; Solomon et al., 2004), although truly separating their influence as we do in the model is currently beyond the reach of experimental methods.

Recurrent network dynamics are unlikely to affect responses on time scales beyond several seconds. However, experimental paradigms that quantify long-term adaptation typically briefly repeat the adapter just before the test (top-up design). As a consequence, such measurements of long-term adaptation (i.e. plasticity) may be confounded with the more short-lived effects reported here (dynamics). This may offer an explanation for some of the contradictory reports found in the fMRI adaptation literature (Krekelberg et al., 2006b). More importantly, removing this confound in future studies requires the introduction of a substantial temporal delay between adapter and test presentation.

Our findings highlight the general need to consider the complex dynamics of recurrent networks to better understand how the brain computes. For instance, the potential functional benefits of adaptation, such as enhanced discriminability of successive stimuli, are commonly ascribed to changes in intrinsic neuronal properties or the strength of synaptic connections (Muller et al., 1999; Krekelberg et al., 2006a; Kristjansson, 2011). Our model, however, shows that a recurrently connected network can generate the same behavior and therefore the same purported functional benefits, without plasticity. While the results reported here focus on visual cortex, analogous dynamic changes in tuning and integration are likely to affect computations throughout the recurrent networks of the brain.

Chapter 3. Neural population dynamics change perceived orientation

3.1 Introduction

The use of sensory adaptation as a tool to study the underlying mechanisms of sensory processing itself has been widely exploited (Hegde, 2009). In particular, the effect of orientation selective adaptation on the perception of subsequent orientations has been of specific interest (Gibson and Radner, 1937; Magnussen and Johnsen, 1986; Harris and Calvert, 1989; for a review see Schwartz et al., 2007). The typical result of tilt aftereffect (TAE) experiments is a repulsive shift on the perceived orientation of oriented gratings away from the adapting orientation for similar orientations, and sometimes a smaller attractive shift for near orthogonal orientations (Gibson and Radner, 1937; Morant and

Harris, 1965; Muir and Over, 1970; Mitchell and Muir, 1976). These findings have been replicated over a variety of timescales (Wolfe, 1984; Wenderoth and van der Zwan, 1989). Our model presented in Adaptation without plasticity, however, predicts an attractive shift in orientation perception for short test durations, for all adapter orientations other than same or orthogonal.

In this chapter we report the results of two experiments that tested for the presence of these attractive shifts in perceived orientation for short test stimuli. In both cases, this was achieved by having participants compare the perceived orientation of a test orientation on one side of the display to that of a "reference" orientation on the opposite side. Crucially, the test orientation was preceded by one of two adapter orientations, while the reference was preceded by a null (i.e. non-oriented) adapter (see Figure 10a and Figure 12a for example trials). Given that the effect predicted by our model is due to slow population dynamics, whereas the TAE is thought to be due to plasticity, we tried to minimize the experimental variables that could induce longer-term plasticity. This was achieved by randomizing the orientation of the reference and test stimulus from trial to trial so that no one orientation was presented more often than others.

3.2 Methods

3.2.1 Participants

Three subjects (two male, one female) participated in Experiment 1, one of whom was naïve to the purpose of the experiment. Eight naïve subjects (three male, five female) participated in Experiment 2. All participants had normal or corrected-to-normal vision, were aged between 17 and 36 and were right-handed. All experimental procedures were approved by the local Institutional Review Board and followed the National Institute of Health's guidelines for the ethical treatment of human subjects. All subjects provided written informed consent.

3.2.2 Apparatus

Stimuli appeared on a Sony FD Trinitron (GDM-C520) CRT monitor using custom software (Neurostim, <u>http://neurostim.sourceforge.net</u>). The display measured 40° (width) by 30° (height) at a viewing distance of 57 cm.

Eye-position was monitored using a video-based eye tracker (500Hz; SR Research, Mississauga, Canada) and the subject's head was stabilized using a bite bar. Monitor resolution and refresh rate were 1024x768 and 120Hz in Experiment 1, and 1280x960 and 90Hz in Experiment 2.

3.2.3 Visual Stimuli

The visual stimuli consisted of oriented Gabors (sinusoidal gratings modulated by a Gaussian envelope to make it fade smoothly into the background) and null-adapters that we constructed by adding together and normalizing eight Gabors whose orientations span orientation space evenly with a random offset. These null-adapters were matched in spatial frequency and contrast to the adapting and testing Gabors, but contain little if any orientation specific energy (see Figure 10 and Figure 12). Null-adapters were also used in our paradigm as backward masks. The central fixation stimulus was a small red square located in the center of the screen, which remained visible for the duration of the trial.

The mean luminance of the equal energy white screen was set to 30 cd/m^2 . For all stimuli (Gabors and null-adapters) the width (i.e. standard deviation) of the Gaussian envelope was 0.8° of visual angle. Gabors had 75% contrast and the gratings in the masks 90% contrast, which produced the most similar perceived contrasts.

3.2.3.1 *Experiment* 1

The centers of the stimuli were three degrees of visual angle to either side of the fixation dot. The spatial frequency of the sinusoidal modulations underlying both Gabors and null-adapters was 2 cycles/°.

3.2.3.2 *Experiment* 2

The centers of the stimuli were twelve degrees of visual angle to either side of the fixation dot. Experiment 2 relied on the inclusion of an adapter of imperceptible orientation, therefore the spatial frequency of the sinusoidal modulation of the adapter was varied per subject, according to their individual ability to perceive orientation at high spatial frequencies. All other parameters in this experiment matched those of Experiment 1.

For each subject in Experiment 2 we ran a preliminary experiment to map their ability to detect the orientation of the Gabors as a function of spatial frequency. This experiment matched experiment 2 except that both test and reference Gabors were preceded and followed by null-adapters. The spatial frequencies for the reference and test Gabors was varied factorially between 3 cycles/degree and 8 cycles/degree. The orientation offset between the test and the reference was set to an easily distinguishable 20 degrees. The

subjects' task was to identify the most clockwise stimulus. From their responses we determined two spatial frequencies per subject. First, the highest spatial frequency for which subjects performed the task at least 80% correct. In the main experiment, the test stimulus was set to this spatial frequency. Second, we determined the lowest spatial frequency for which the subject was at chance performance. In the main experiment, the adapter stimulus was set to this spatial frequency.

These initial experiments resulted in the following spatial frequency settings in the main experiment: for 5 out of 8 subjects the spatial frequency of the test and null-adaptors was 4 cycles/°, for the remaining 3 subjects the spatial frequency was 6 cycles/°. The adapter grating had a spatial frequency of 6 cycles/° (1 subject), 7 cycles/° (4 subjects), or 8 cycles/° (3 subjects).

3.2.4 Procedure

Subjects were required to maintain fixation within a 3° by 3° square at the center of the display (around the fixation point) for the duration of each trial, without including the response epochs. Each trial started when the subject fixated; trials in which subjects failed to fixate appropriately were terminated immediately, discarded, and repeated randomly at a later time within the block.

After a delay of 250 ms there was a simultaneous presentation of an oriented adapter (a Gabor) on the left or right (selected randomly) and a null adapter (as described in Visual Stimuli) on the other side. The null adapter was included in this instance of the experiment to match the spatial frequency and contrast adaptation with respect to the

orientation adapted side and to maintain relative symmetry in the display (to prevent shifts of attention to one side). Immediately after, two target stimuli were presented (oriented Gabors), one on each side of the fixation point for a variable duration of 50, 100 or 200 ms. Subjects were instructed to determine which one of these two oriented Gabors was tilted more clockwise. The Gabor that appeared in the same location as the adapter is referred to as the "test" stimulus (right Gabor in the example trial shown in Figure 10a), and the Gabor that appeared in the location of the null-adapter is the "reference" stimulus (left Gabor in the example trial shown in Figure 10a). After the target stimuli, two nulladapters were presented in the same spatial locations on either side of the screen for 500 ms. These masks served to minimize afterimages produced by the test and reference stimuli and to limit the amount of temporal integration of the target stimuli within the visual system. Subjects indicated their response by pressing one of two designated keys on the keyboard after the test stimuli disappeared.

In a traditional adaptation paradigm, the same (or similar) adapter orientation is typically repeated on every trial with the intent of accumulating adaptation (i.e. plasticity) effects across trials. Here, however, we wanted to minimize such effects to isolate the recurrent network dynamics. Therefore, we randomly chose the orientation of the reference on each trial to prevent build-up of adaptation and presented the test at a range of orientation offsets relative to the reference.

3.2.4.1 Experiment 1

The orientation of the adapter was set to be 20° either clockwise or counterclockwise relative to the orientation of the reference. The orientation of the test was offset from the

orientation of the reference by -12, -8, -4, 0, 4, 8, or 12°. The adapting stimulus was presented for 100 ms.

3.2.4.2 Experiment 2

We used an adaptive staircase procedure (Kontsevich and Tyler, 1999) to choose the orientation offset between the test and reference Gabors in each trial based on the previous responses. This method (as opposed to the method of constant stimuli) was used to avoid making a-priori assumptions about the sensitivity of each subject for this task. The adapting stimulus was presented for 200 ms.

We randomly interleaved catch trials (12.5% of trials) to test whether subjects were, contrary to instructions, responding directly to the adapter (and not the test stimulus). In these trials the test Gabor was replaced by a null-adapter. Given that in these trials the test had no orientation, the only oriented information on that side of the screen was that of the adapter. If the subject was able to perceive the orientation of the adapter (even though it was previously adjusted to be outside the visible range; see above), we would expect its orientation to bias the subject's responses. Accordingly, we quantified the proportion of "test clockwise" responses as a function of adapter orientation for catch trials as an additional measure of adapter visibility during the main experiment.

3.2.5 Trial presentation

Trials were presented in blocks of 64 trials for each test duration. Each subject completed a minimum of 60 blocks in Experiment 1 and 96 blocks in Experiment 2. All subjects received between two and five hours of training on different versions of the task prior to completing experimental trials. Data collected during the training blocks were not analyzed.

3.2.6 Data analysis

3.2.6.1 Experiment 1

We fitted logistic psychometric functions to each subject's responses using the *psignifit* toolbox version 2.5.6 (Wichmann and Hill, 2001b, a) in MATLAB version 8.2.0.701. These functions relate the orientation difference between the reference and test to the likelihood of responding "test clockwise" (e.g. Figure 10b). Accordingly, the inflection point corresponds to the offset required for the reference and test to be perceived as the same orientation (i.e. the point of subjective equality; PSE). We first fitted a psychometric function to the data pooling across adapter orientation to estimate the lapse rates for each subject for each adapter condition. We later fitted each curve separately, fixing the lapse rates to the previously calculated values. We determined a p value for the PSE difference between the curves using the function *pfcmp* (two-tailed) (Wichmann and Hill, 2001b, a). We defined "shift" as the difference in the PSE's of the adapter -20 degrees condition and the adapter 20 degrees condition. A positive difference signals an attractive shift.

3.2.6.2 *Experiment* 2

For the regular (non-catch) experimental trials we followed the same analysis as for Experiment 1. In the catch trials, the subjects could only compare the high-frequency adapter to the low-frequency reference (no other oriented patterns were on the screen in these trials). We tested whether the subject responses were influenced by the adapter orientation under these conditions. For this we calculated the proportion of 'test clockwise' responses for each adapter orientation, and tested whether these proportions were significantly different using a two-tailed Chi-square test, *chi2stat* in MATLAB version 8.2.0.701.

3.3 Results

We designed a behavioral task to measure whether at short timescales human observers exhibit a perceptual shift consistent with the predictions of the model presented in Chapter 2. To relate such neural changes to perceptual consequences, we assumed that V1 neurons are labeled lines for orientation. Under this assumption, common in sensory decoding studies, a spike from a neuron is a vote in favor of that neuron's preferred orientation. Given this assumption, the attractive shifts in the population response of the model (Figure 5) predict an attractive shift in perception – that is, a shift toward the adapter orientation. Such a perceptual shift is opposite in direction to the standard repulsive tilt after-effect (Gibson and Radner, 1937).

3.3.1 Experiment 1

Subjects were instructed to respond which of two oriented Gabors (presented on each side of the fixation stimulus) was tilted more clockwise (Figure 10a shows an example trial). One of these Gabors (the 'test') was preceded by an oriented adapter at the same spatial location, whereas the other ('reference') was preceded by a null-adapter stimulus composed of eight uniformly distributed orientations. The 'test' screen was presented for



Figure 10 - Psychophysical measurement of short-term adaptation effects in humans. **a:** Example trial. Subjects fixated centrally for the duration of the trial. An oriented adapter and null-adapter appeared on either side of the fixation stimulus for 100 ms. Immediately after, an oriented test stimulus appeared on the same side as the adapter while a reference stimulus appeared on the opposite side, followed by null-adapters in both locations. At the end of the trial, subjects responded which of the oriented Gabors in the test display (left or right) was tilted more clockwise. **b:** Psychometric curves for one example subject, for three test durations: 50, 100 and 200 ms. For the shortest duration of the test stimulus (50 ms), the perceived test orientation was significantly shifted towards the orientation of

the adapter. This demonstrates an attractive shift in orientation perception. Box plots show PSE confidence intervals (percentiles at 2.3, 15.9, 84.1, and 97.7).

The model predicts that perception of the test orientation should be biased towards the adapter orientation, but only at shorter durations since continued stimulation would lead to convergence of the population response, i.e. disappearance of the shift in tuning curves. The results confirmed this prediction. Figure 10b shows psychometric curves for one representative subject. The shortest test duration shows an attractive shift in orientation perception, which in this data format corresponds to a shift of the PSE *away* from the adapter. Indeed, the point of subjective equality for the clockwise adapter (green curve) was shifted counter-clockwise compared to that associated with the counter-clockwise adapter (blue curve). The three subjects tested all had attractive shifts at the shortest test duration (p<0.05 using *pfcmp* from *pfit* toolbox (Wichmann and Hill, 2001b, a)), but the shifts disappeared for longer test durations (Figure 11).

Experiment 1 confirmed an attractive bias in perception towards the adapting orientation for short durations of the test stimulus. However, an alternative explanation for our results could be that subjects mistakenly compared the adapter (rather than the test) to the reference orientation. If subjects always did this, then they should respond "clockwise" and "counter-clockwise" on every trial in the clock-wise and counter-clockwise conditions, respectively (i.e. both psychometric curves should be flat and have a vertical offset between them). Results of this nature were observed in a comparable experiment designed to test the effects of saccades on perceived orientation (Paradiso et al., 2012). We did not observe such extreme biases in our experiment, but even if subjects responded to the test most of the times, and to the adapter on only some trials, an attractive bias would be expected.



Figure 11 - Difference in PSEs between 20 and -20° adapting conditions, as a function of test duration. For subjects 2 (pink) and 3 (green), the only significant difference is for 50 ms test duration (p<0.004). For subject 1 (blue), who was not tested at the 50 ms test duration (due to poor performance), the difference is significant for 100 ms test duration (p<0.0001) and not significant for 200 ms test duration.

3.3.2 Experiment 2

To address this limitation we designed a second experiment in which we exploited the phenomenon that a high spatial frequency grating adapter can induce a perceptual aftereffect measured with a low spatial frequency test stimulus. Importantly, this crossfrequency adaptation occurs even when the frequency of the adapter grating is too high to be resolved (He and MacLeod, 2001; Rajimehr, 2004). Therefore, if the adapter orientation cannot be seen, it cannot bias the results in the way that it could have in Experiment 1. Nevertheless, assuming some coupling between the spatial frequency channels that encode the adapter and the test spatial frequencies (He and MacLeod, 2001), a perceptual shift could still be observed.



Figure 12 - Task and group data for Experiment 2. **a:** Example trial. Experiment 2 is similar to experiment 1, except that the spatial frequency of the underlying Gabors in the adapting interval was too high for subjects to discriminate their orientation. The adapters were presented for 200 ms. **b:** Perceptual shifts for the 8 subjects, for each test duration (i.e. the difference in their PSEs for clockwise and counter-clockwise adaptation). Error bars are 95% confidence intervals from the Monte Carlo simulations in pfcmp.

In brief, we increased the spatial frequency of the adapting Gabor and null-adapter, and moved them to a more peripheral location, such that subjects were not able to discriminate their orientation (Burr and Wijesundra, 1991) (Figure 12a shows a cartoon of an example trial). The non-resolvable spatial frequency was determined per subject in preliminary sessions and we used catch trials in the main experiment to confirm that subjects were not responding to the adapter orientation (Chi-square test, $p \ge 0.67$ for all test durations and subjects; see Methods).

Nevertheless, as in the first experiment, we again found shifts in the psychometric functions that were consistent with an attractive perceptual after-effect. The individual values for the PSE shifts (Figure 12b) were variable, but a repeated measures one-way ANOVA revealed a significant effect of test duration on the mean shifts across the group (F(2,14) = 37.79; p<0.001). Post-hoc comparisons (Bonferroni corrected; Figure 13) showed that the attractive shift for the short duration was significantly greater than the shifts for the intermediate and long durations (p<0.05), which were not significantly different from each other.



Figure 13 - Mean shift across subjects as a function of the duration of the test stimulus. Error bars = 95% confidence intervals. There was a significant attractive shift for short test stimulus durations that disappeared for longer durations. This confirms the attractive shift in orientation perception on short time scales.

3.4 Discussion

We found and confirmed a novel time course of adaptation that generates an attractive shift in the perceived orientation of gratings relative to the adapting orientation, which is opposite in sign to the better known TAE. Our behavioral findings are consistent with the model predictions presented in Chapter 2. Because the model predictions depend strongly on the same population of neurons responding to both the test and adapter, our finding furthermore predicts that the functional connectivity between orientation sensitive neurons is relatively independent of their spatial frequency preference.

There are several reasons why other groups probing the system at similar timescales may not have found this effect: first, the magnitude of our effect is relatively small compared to the TAE, and if care is not taken to minimize the TAE repulsive effects, these may swamp the smaller attractive shifts due to population dynamics. For example, most TAE experiments were performed using a near vertical adapting stimulus for the convenience of instructing subjects to report clockwise or counterclockwise perceived orientations (relative to vertical) for the test (Gibson and Radner, 1937; Morant and Harris, 1965; Muir and Over, 1970; Campbell and Maffei, 1971; Mitchell and Muir, 1976; Wenderoth and van der Zwan, 1989). The repetition of a similar adapting orientation over every trial optimizes the system to accumulate effects of plasticity, whereas we used uniformly distributed adapting orientations across trials. Second, many groups have observed evidence to support the existence of an "oblique effect", which refers to a preference of the visual system towards vertical and horizontal orientations (Mach, 1861; Campbell and Maffei, 1970; Maffei and Campbell, 1970). If the mechanisms underlying the repulsive TAE are more sensitive to cardinal orientations as well, then the near vertical orientations used for previous TAE experiments could exacerbate the repulsive effects, further hiding the population dynamics effect.

New timescales continue to be found for orientation specific aftereffects, even for up to several seconds (Fischer and Whitney, 2014). However, given that humans, as many other mammals, make several eye movements per second during normal exploratory behavior (Carpenter, 1988), the timescale of a few hundred milliseconds is highly relevant for behavior. In addition, several potential functional benefits of adaptation have been proposed, such as enhanced discriminability of successive stimuli (Muller et al., 1999; Krekelberg et al., 2006a; Kristjansson, 2011). The attractive shift in perceived orientation found in this study not only provides evidence for lingering activity in cortical

primary visual cortex, such as aiding integration of sensory inputs over short timescales.

Chapter 4. Dynamic response properties in primate V1

4.1 Introduction

In Chapter 2 we showed that the model population response and tuning properties depended on the recent history of stimulus orientations. Further, these history effects depended critically on the weight profile of the intra-cortical recurrent connections (e.g. larger and more persistent shifts for stronger intra-cortical connections). These results suggest that recurrent connectivity plays a key role in shaping how V1 neurons respond to rapidly changing stimulus conditions. Inspired by this observation, we recorded the responses of V1 neurons to rapid sequences of oriented stimuli in order to gain insights into the contributions of recurrent inputs to cortical representations of orientation. Felsen et al. (2002) used this kind of stimulus on anesthetized cats, but only reported average tuning curve effects. Benucci et al. (2009) also used similar stimuli, reporting more on the dynamics of the average tuning curve. However, both of these studies were performed on cats under anesthesia, which could have profound effects on network-related phenomena.

To our knowledge these types of experiments have not yet been performed in awake behaving primates. In addition, we wanted to have full control on the experimental details and data, as well as analysis, to directly compare the neural behavior to the behavior of the recurrent model. Therefore, we recorded from primary visual cortex of three awake behaving macaques chronically implanted with multi-electrode arrays. The animals performed a fixation task while being presented with sequences of oriented gratings. We report a relatively raw and unprocessed measure of neural activity: the multi-unit activity envelope (MUAe), which is thought to reflect the activation of multiple neurons that are located 100-200 µm around the electrode tip (Super and Roelfsema, 2005). We investigated whether the changing response properties of the V1 MUAe to dynamic stimulus sequences can be explained fully in terms of a recurrent model (as suggested by the analysis of Chapter 2) or instead (or in addition) requires plasticity in the lateral connectivity (as argued by Felsen et al., 2002). Further, we assessed whether the dynamics of V1 responses in the grating stream paradigm contain significant nonlinearity or instead can be predicted fully from the linear summation of the responses to individual stimuli (as argued by Benucci et al., 2009).

4.2 Methods

4.2.1 Electrophysiology

Three adult male macaque monkeys (*Macaca mulatta*) were implanted with 32-channel PlextrodeTM Floating Microelectrode Arrays (FMA) in parafoveal primary visual cortex (V1) under sterile conditions and using isoflurane anesthesia. Electrodes were platinum/iridium, for further details visit the <u>website</u>. The length of the electrodes was between 0.6 and 1.5 mm and the nominal impedances of the electrodes at the time of insertion was between 0.7 and 1 M Ω . One of the monkeys (M) was implanted with one array in the right-hemisphere with receptive fields inside the square [-1 1]x[-4 -2] (in degrees of visual angle); another (Y) with two arrays in the left-hemisphere with receptive fields in [-3 -1]x[-3 -1] and [-4 -2]x[-4 -2]; and the third (N) with two arrays on the right hemisphere with receptive fields in [-4 -2]x[-6 -4] and [-4 -2]x[-4 -2]. During the

experiments, the animals sat in a primate chair with their heads restrained and received liquid rewards for each completed trial. Experimental and surgical protocols were approved by the Rutgers University Animal Care and Use Committee. The protocols were in agreement with National Institutes of Health guidelines for the humane care and use of laboratory animals.

Following standard procedures in the laboratory, the electrical signals were amplified, filtered, digitized (25 kHz) and stored using an Alpha Lab system (Alpha-Omega Engineering), or at 30 kHz using a Trellis recording system (Ripple Inc). Eye position was recorded at up to 2 kHz using an infrared eye tracker (Eyelink).

To obtain the multi-unit activity envelope for each channel, we rectified the continuous signal, band-pass filtered the result between 750 and 5000 Hz using a Butterworth filter, and down-sampled the signal to 1250Hz (Super and Roelfsema, 2005).

4.2.2 Visual stimuli

Visual stimuli were displayed on a 20-inch CRT monitor (Sony GDM-520) at 57 cm distance using custom software called Neurostim (http://neurostim.sourceforge.net). The display covered 40° (width) by 30° (height) of the visual field and had a resolution of 1024 x 768 pixels. Except for a red fixation point, all stimuli were presented in equal energy white (x = 1/3, y=1/3, Commission Internationale de l'Eclairage color space) with luminance ranging from 0.2 to 85 cd/m2, and the mean luminance of the screen was 30 cd/m2.

The stimulus consisted of full screen gratings of 50% contrast flashed in a pseudorandom sequence for 5 frames each at a frame-rate of 150 Hz (i.e. 33.3 ms for each orientation; see Figure 14 for an extract of an example trial). We ran pilot experiments where we recorded neural response to a range of spatial frequencies and found that on average 1 cycle/degree evoked the largest responses across the array channels. Therefore the spatial frequency of the gratings was set to 1 cycle/degree. Blank screens of the same mean luminance and duration, but of uniform gray color, were interleaved with the gratings with a 30% probability.

4.2.3 Procedure

A trial started with the presentation of a central red fixation target. The animal was required to fixate within a 3° by 3° window surrounding the fixation dot until the end of the trial (3.5 s, equivalent to 105 orientations per trial) to receive a juice reward. In practice, these well trained animals typically fixate very tightly on the fixation dot regardless of the fixation window size. Trials in which fixation was not maintained were immediately aborted and the data obtained during fixation were only analyzed if the animal fixated for at least 500 ms.

4.2.4 Data analysis

4.2.4.1 Tuning curves and principal components analysis (PCA)

The neural response latency was estimated to be around 60 ms from the event related responses (Figure 19). The response to an oriented grating was defined as the mean response for the 33 ms after grating onset plus latency. Tuning curves were obtained by

collating the responses to each of the unique orientations across all trials and then calculating the average MUAe responses. To select the tuned electrodes, we fit von Mises curves to the tuning curves (see equation (3) for an example), and used Akaike information criterion (Burnham and Anderson, 2002) to decide whether the von Mises fit the tuning curve better than the mean, despite the additional parameters (i.e. whether the electrode was significantly tuned). If the electrode was tuned, we then subdivided the pooled responses to each orientation according to the preceding, or "adapting", orientation. This resulted in 12 tuning curves, one for each adapting orientation, and for each we repeated the same fitting procedure. Based on previous findings (Felsen et al., 2002) and our modeling work (Chapter 2), we predicted that tuning curves would shift in a manner that depended on the adapting orientation. We defined the shift for each adapting orientation to be the difference between the peak of the von Mises fit for that adapter tuning curve and that of the global, pooled tuning curve. For this analysis we only considered the tuned MUAes, i.e. those with 12 defined shifts.

We anticipated that the pattern of shifts across all possible adapter orientations could be qualitatively different across electrodes (e.g. some repulsive, some attractive, etc.). To unpack this variability, we performed a principle components analysis on the shift measurements. We first normalized the shifts of each MUAe to the maximum shift, so that MUAes with larger shifts wouldn't dominate the analysis. We then performed the PCA using MATLAB's *pca* function.

4.2.4.2 Event-related analysis

We first subtracted the mean across trials from the MUAe signal for each electrode to remove the onset response to each grating that occurred regardless of the grating sequence. We took the responses to each oriented grating 50 ms before to 250 ms after its onset, and z-scored them by dividing the mean across repeats by the standard deviation across repeats. These event related responses were then subdivided into smaller groups depending on whether the preceding orientation was the same orientation (0°), 90, 45, or -45° away; and another group where each orientation at 0 ms was followed by a uniform gray screen at 33 ms. To simplify the comparison across recording sites (and allow averaging across sites), we relabeled the stimulus orientations in terms of their angle relative to the preferred orientation (i.e. as if all multi-units preferred 0°). The mean event related response for a given recording session (i.e. day) was calculated as the average over all simultaneously recorded electrodes (Figure 19, Figure 21a and Figure 23a).

A key aim was to compare these orientation-pair responses with those expected based on a linear summation of the component grating responses. Predicted responses were generated by taking the average response to each grating (irrespective of preceding orientation) and then summing them, appropriately time-shifted (i.e. with the second component in the pair shifted backward by 33ms). The difference between the predicted and actual pair responses (residuals) was computed to provide insights into the nonlinearity in the representation of orientation.

4.3 Results

We recorded from primary visual cortex using chronically implanted arrays. Our stimuli consisted of 3.5 second streams of gratings in which each orientation was presented for 33 ms each (for an example trial, see Figure 14). Interleaved with the oriented grating intervals (with a 30% probability) were blank intervals consisting of a uniform gray screen with same mean luminance as the gratings (30 cd/m2). We investigated the population activity of the multi-unit activity envelope (MUAe) (see Methods for a description of how this measure was obtained). We analyzed the activity only for electrodes that had well-tuned MUAe (see Methods for selection criteria).

The MUAe measures the action potentials that are near the electrode tip, which works especially well when the neurons near the tip have similar response properties (Super and Roelfsema, 2005). Because there is no thresholding involved (as opposed to both single cell recordings and regular MUA signal), it does not have the problems of bias towards larger neurons. Instead, the MUAe provides a continuous measure of all high-frequency activity (up to 5000Hz). Tuning, latency and receptive field properties of the MUAe were found to be comparable to that of single unit activity (Lamme and Roelfsema, 2005).



Figure 14 - Extract of an example trial from experiment. The gratings were 50% contrast, with a spatial frequency of 1 cycle/°, had a randomly chosen spatial phase and one of 12 orientations. Randomly interleaved were blank frames of the same mean luminance (30 cd/m2) but 0% contrast, which occurred with a probability of 30%.

We analyzed the stimulus history effect on the orientation tuning of MUAe. We first obtained an average orientation tuning curve for each MUAe, by averaging the response to each orientation in the stream independently of the preceding or following orientations. By fitting a von Mises curve, we were able to estimate the preferred orientation for each MUAe. We repeated this analysis by subdividing the responses according to each preceding orientation, and obtained 12 tuning curves for each MUAe, one for each "adapting" orientation. We then assessed the difference in preferred orientation for each of these curves relative to the average tuning curve, obtaining a shift for each adapting orientation, as in Figure 7 a and b (Figure 15). For consistency, Figure 15, Figure 16, and Figure 19 to Figure 23 correspond to data obtained from M in a single recording session, but qualitatively similar results were observed for all 7 recordings for each of the three animals that we have analyzed so far.



Figure 15 – Shifts in the MUAe tuning curves for 30 array electrodes recorded in one session (electrodes 9 and 26 did not have significant tuning, and were therefore excluded from the analysis). White bars correspond to significant tuning curves for each adapting orientation, and black bars to adapting orientations for which the tuning curves were not significantly tuned.

For the MUAe we typically observed repulsive shifts, mostly symmetric around the preferred orientation, albeit with varying amplitudes (Figure 15). To study these shifts more systematically, we performed a principal components analysis. For each recording

session, we first set the preferred orientation to be zero and then normalized the shifts to be between -1 and 1, so that the PCA components would not be biased to units with larger shifts (Figure 16 a).

The first three principal components (Figure 16 b) accounted together for more than 95% of the variance in the data (Figure 16 c). The first component alone (blue curve in Figure 16 b), which captured the repulsive symmetric shift in preferred orientation as a function of adapting orientation, accounted for over 80% of the variance. The second component (red curve in Figure 16 b), served to introduce a small amount of asymmetry in the shifts, and we suspect the third (orange curve in Figure 16 b), reflects the influence of slight errors in the estimation of preferred orientation of each recording site. The projection scores of the data onto the first principal component (Figure 16 d) show that the shifts were always repulsive (all positive PC 1 projection scores).


Figure 16 - Principal components analysis for the multi-unit activity envelope (MUAe) of 30 simultaneously recorded electrodes with significantly tuned tuning curves. **a:** Shifts for all MUAe, normalized between -1 and 1. **b:** Three first principal components for the shifts presented in **a. c:** Percentage of explained variance for each of the principal components (only shown for components 1 to 6). **d:** Projection scores of the first two principal components, for the 30 curves shown in **a.**

We simulated the same experiment on our C-model and M-model from Chapter 2 (Cmodel shown in Figure 17, but the M-model produces nearly identical results) and performed the same set of analyses. The similarities are striking, from the normalized shifts, to the first principal component, and even the fraction of variance explained for the components.



Figure 17 - Principal components analysis for 18 C-model units with equally spaced preferred orientations. **a:** Orientation tuning curve shifts of C-model units, normalized between -1 and 1. **b:** Three first principal components for the shifts presented in **a. c:** Percentage of explained variance for each of the 6 first principal components. **d:** Projection scores for the first two principal components, for the 18 curves shown in **a.**

We were interested in understanding the MUAe responses further: exploring the time course of the responses, rather than only looking at static tuning curves; studying the responses to pairs of stimuli with a fixed orientation difference; and how this response to a pair related to the sum of the responses to the individual orientations. We performed an event related analysis by averaging a window of the response time locked to the presentation of a certain orientation, as explained in detail in Figure 18.



Figure 18 - Obtaining the event related average response of one electrode. **a:** A section of a trial. The upper panel shows the stream of gratings presented on the screen, and the lower panel (blue trace) the MUAe recorded from one electrode. The three colored rectangles delineate 300 ms windows around each occurrence of the orientation 120°, marked by the dashed vertical lines inside each rectangle. **b:** The response corresponding

to each of the windows is aligned to the onset of the relevant orientation, averaged, and divided by the standard deviation to get a z-scored response to each orientation. **c:** The tuning curve for the electrode is calculated to express stimulus orientation relative to preferred. **d:** The z-scored response obtained in **b** is assigned to its appropriate orientation relative to preferred, and the z-score modulation is expressed using color. **c:**This process (**a** - **d**) is repeated for each stimulus orientation to obtain the event related average response of each electrode.

Figure 19a shows the dynamics of the response separately for each orientation, averaged over all tuned electrodes in one session. Data are plotted as z-scores (mean response across repeats divided by the standard deviation across repeats for each electrode). Thus, a horizontal slice at 0° shows the response to the presentation of the preferred orientation. The latency of the response was about 60 ms, and the response lasts for around the duration of that oriented stimulus, i.e. ~33 ms. Each vertical slice across Figure 19a can be interpreted as the tuning curve of the MUAes (at a certain point in time after stimulus onset). The C-model and M-model (Figure 19 b and c) have different amplitudes of response, but have a qualitatively similar single orientation response. Note that in Chapter 2, the models received input from LGN instantly; i.e. there was no transmission latency. To aid comparison with the experimental data, here we assumed a delay of 40ms for the first inputs to the model V1.



Figure 19 - Event related average response to each orientation relative to preferred for 30 well-tuned simultaneously recorded electrodes. The response (for each electrode in **a**, and for each model unit in **b** and **c**) was z-scored as the mean response across repeats divided by the standard deviation across repeats. **a:** MUAe. **b:** C-model. **c:** M-model.

To visualize the interaction between successive oriented stimuli (pairs), we extracted four subsets of the stream of gratings: pairs in when an orientation was preceded by the same orientation (0° jump, Figure 20), the orthogonal orientation (90° jump, Figure 20) or an orientation 45° away in either direction (45° and -45° jumps, Figure 20). The y axis in Figure 20 represents orientation of the *second* stimulus in the pair relative to the preferred orientation (hence the increased activity at 0° at the time of the second grating in all

plots), and different jump sizes are shown in different panels. In this format, each horizontal slice through a heat map corresponds to the dynamics of the MUAe response for a different stimulus pair but with a fixed jump size. For example, a horizontal slice at 0° in the 90° jump panel in Figure 20a represents the average z-scored MUAe response to the anti-preferred direction presented at time -33 ms (i.e. 33 ms before the onset of the second grating) followed by the preferred direction at 0 ms; whereas a horizontal slice at 90° represents the average z-scored MUAe response to the preferred orientation presented at -33 ms followed by the anti-preferred orientation at 0 ms. Note that the response latency is around 60 ms, therefore the first bump in activity occurs between 30 and 60 ms, and the second bump between 60 and 90 ms. It should be further noted that vertical slices in these graphs represent tuning curves (as in Figure 19), not population responses.

The tuning curves reveal some interesting non-linear interactions between successive stimuli. For instance, the response to the second stimulus in the 0-jump pair is reduced, while the response to the second stimulus in the 90-jump pair is enhanced. The \pm 45-jump pairs show the smooth transition of the tuning curves which results in tuning curve shifts as shown in Figure 16 above.

For comparison, we subjected the C and M-models to the same analysis (Figure 20 b and c). Neither the C nor the M-model captured the nonlinear suppression in the response to the second stimulus for a 0° jump (i.e. repeated presentations of the same stimulus). This suggests that a mechanism other than recurrent connectivity may be needed to explain this effect (see Discussion). The nonlinear enhancement in the MUAe response to the

second stimulus in the case of a 90° jump, however, was also present in the C-model, indicating that it can be understood as an effect of recurrent connections.



Figure 20 - Event related responses to jumps in orientation, aligned to the onset of the second grating in the pair (at 0 ms, response seen after ~60 ms latency). **a:** Event related response to a 0, 90, 45 and -45° jump in orientation, respectively. MUAe activity is plotted with respect to the orientation of the second grating relative to preferred orientation (y-axis). Color map is the same as in **Figure 19a**. **b:** Same as **a** for the C-model. Color map is the same as in **Figure 19b**. **c:** Same as **a** for the M-model. Color map is the same as in **Figure 19b**. **c:** Same as **a** for the M-model. Color map is the same as in **Figure 19c**.

Benucci et al (2009) have claimed that the responses to two successive orientation stimuli can be understood as the linear sum of the responses to the individual stimuli. Such linear behavior would contradict the main conceptual message of the recurrent network, i.e. that reverberations in activity influence the dynamic responses to stimuli presented in fast succession (which would be expected to have non-linear effects). Hence, we tested the linear prediction explicitly by constructing the predicted response (Figure 21a) based on the linear summation of two event related average MUAe responses (Figure 19a). To enable a direct visual comparison of the linear prediction with the actual response to paired gratings, we also calculated the difference between the measured paired response (Figure 20a) and the linearly predicted response (residuals: Figure 21b).



Figure 21 - Summing two responses to a single orientation does not explain the response to jumps in orientation – one recording session. **a:** Predicted event related response to 0, 90, 45 and -45° jumps in orientation if the response was just the sum of two single (time shifted) orientation responses **b:** Difference between MUAe and linearly summed responses (**Figure 19a** - **a**). The clearest systematic failures of the linear model were in the 0 and 90° jumps. Color-map values are the same as in **Figure 19a**.

Figure 21a shows that the qualitative predictions of the linear model are indeed similar to the observed responses to paired gratings (Figure 21a). The residuals (Figure 21b), however, show that the linear model (just like the C- and M-model) failed to capture the suppressed response to the second stimulus for the 0° jumps (same orientation presented

twice). The linear model also failed to capture the enhancement in the response found after 90° jumps. Finally, the linear model failed to capture the 45 and -45° jump responses, as evidenced by the orientation specific residuals (though these will become more apparent in our next analysis, below).

We averaged these residuals over 7 recording sessions for monkey M, after which their structure became much more apparent (Figure 22). There is a generalized suppression to the second stimulus for the 0° jumps, but the suppression is orientation specific, i.e. it is stronger for $0^{\circ}-0^{\circ}$ (preferred orientation followed by preferred orientation) than for 90° -90° (anti-preferred orientation followed by anti-preferred orientation). In addition this suppression is followed by a very brief enhancement, especially near the preferred orientation. The enhancement in the response after 90° jumps is strongest for 90° - 0° and $0^{\circ}-90^{\circ}$ (jumps from preferred to anti-preferred orientation and vice-versa). This enhancement is followed by a longer lasting, but smaller in magnitude, suppression, that does not appear to be orientation specific. The residuals for the $\pm 45^{\circ}$ jumps are orientation specific and present in the response to both stimuli, rather than only the second. We speculate that the residuals to the first stimulus are due to the shifts in the tuning curves showed in Figure 16, which are not present in the linear prediction. In addition, although the residuals for the 0 and 90° jumps were highly consistent across animals, the residuals for the $\pm 45^{\circ}$ jumps were more variable.



Figure 22 - Summing two responses to a single orientation does not explain the response to jumps in orientation – average over 7 recording sessions. Average residuals for each of the 4 analyzed orientation jumps (as in **Figure 21**b) for 7 recording sessions.

As we saw in Chapter 2, cortical connectivity plays a key role in determining the persistence of neural responses over time (for example, see Figure 7). Recurrent reverberations allow neural responses to persist through a blank interval for a time that is orders of magnitude longer than the membrane time constant. Accordingly, it is informative to examine what happened to the event related response when the stimuli were followed by a uniform gray screen. The gray screen elicited a reduction in the response (Figure 23a), which was more evident once we took the mean across all preceding orientations (Figure 23b).

A more interesting effect, however, revealed itself after subtracting this generalized blank suppression from the responses to each grating-blank pair, leaving only the orientation-specific activity decay throughout the blank interval. Figure 23c shows a response persistence into the blank interval, which is most pronounced for the preferred orientation. Both the C and M-models exhibit similar persistence as the MUAe data, suggesting that this persistence could be explained by the recurrent connectivity within the networks.



Figure 23 - Persistence in event related responses when followed by a blank screen. **a:** Event related response of the MUAe, C-model, and M-model to all orientations followed by a grey screen with same mean luminance as all other stimuli, but 0% contrast. Color map is the same as in **Figure 19a**, **b**, and **c** respectively. **b:** Mean response over all orientations in **a**. **c:** Event related response (**a**) minus the mean over all orientations (**b**). Color map is the same as in **Figure 19a**, **b**, and **c** respectively.

4.4 Discussion

We investigated the multi-unit responses to rapid sequences of oriented stimuli in primary visual cortex of awake, behaving macaques. These responses revealed strong interactions between successive stimuli that resulted in shifts of the orientation tuning curves. The magnitude of the shifts was even more pronounced than that reported in anesthetized cat visual cortex (Felsen et al., 2002), always repulsive, and mostly symmetric around the preferred orientation (Figure 16). We also found rapid dynamic modulations of response strength (enhancement and suppression) that have not been reported before.

The shifts in preferred orientation for tuned MUAe (Figure 16) were remarkably similar to those found in the model (Figure 17). This supports our observation in Chapter 2 that they do not require plasticity, and are fully explained by recurrent connectivity within the network.

Beyond a confirmation that the recurrent model is appropriate for the neural responses in the awake, behaving primate, however, the data also revealed further complex neural responses to dynamic sequences of oriented stimuli. These included generalized, as well as orientation specific, suppression and enhancements in the response. These responses highlight the importance of understanding the visual cortex as a dynamical system and serve to constrain the model further.

We did not yet fit the recurrent model to the MUAe data, therefore we were not able to make quantitative comparisons between the recurrent model and the MUAe. To find the parameterizations that capture not only tuning curve shifts, but also the full suppression, enhancement, and persistent dynamics of the V1 MUAe, however, may require further fine-tuning of the optimization procedures, which we will do in future work.

The current work, however, already allowed a qualitative comparison with the recurrent models of Chapter 2: these models captured the enhanced response to the second stimulus

in the case of a 90° jump in orientation (Figure 20), as well as the persistence in the response to near preferred stimuli when followed by a gray screen (Figure 23). These phenomena could not be explained in the linear summation model and emphasize the importance of understanding nonlinear dynamics and the lateral connections that may underlie these phenomena.

Finally, the qualitative models also revealed at least one phenomena that may require an extension of the model. Such phenomena is the suppression observed in the response to the second stimulus in the case of a 0° jump (Figure 20). This suppression was not present in either the C-Model, or the M-model, or the linear summation model. It seems that it requires other mechanisms, such as spike-rate adaptation (Sanchez-Vives et al., 2000), that can explain reduced responses for sustained stimulation, to be fully understood.

In a related study, Benucci et al. (2009) found that the linear prediction model explained well the dynamic neural responses in anesthetized cat V1, except for the persistence in the response to near preferred orientations when followed by a blank screen. Based on the description of their methods, we believe that they subtracted the mean event-related response across all orientations from each of the neural responses. Even though they justify this by saying that these responses, which they term the "untuned response components" are not orientation specific and likely dominated by noise, we found this not to be the case in our recordings. They are not orientation-specific in the sense that they represent the mean across orientations, however, we found that they were very different for the different jump sizes. Subtracting the 'untuned' component, for example, removed

the generalized suppression to the second stimulus we observed in the 0° jumps, and the generalized enhancement in the response to the second stimulus we observed for 90° jumps (Figure 20, Figure 21 and Figure 22). Therefore, this 'untuned' component contains the orientation jump-size specific interaction.

We further note that even if these 'untuned' dynamic responses were deemed uninteresting, and therefore subtracted from the paired responses, we still found orientation specific aspects in the dynamic responses for which the linear prediction failed (Figure 22, described in Results).

Further differences of our data compared to Benucci et al. (2009) may also stem from the type of signal analyzed: whereas they analyzed single- and multi-unit spike traces, low-pass filtered with a cutoff at 25 Hz; we opted to analyze the MUAe, Butterworth filtered between 750 and 5000 Hz. This can explain the differences in the latency of the responses as well. In addition, they used anesthetized cats whereas we recorded form awake behaving macaques. Further, they did not examine patterns in the residuals, but rather analyzed the ability of the linear model to explain the data as the percentage of variance explained of the tuned data (subtracting the 'untuned' component).

The current analysis focused on multi-unit activity, which reflects the aggregate response of the neurons in a 100 - 200 micron radius around the electrode (Super and Roelfsema, 2005). The abstract model developed in Chapter 2 may be most appropriate to describe such aggregate responses, but we are also interested to determine whether the model can

capture the responses of single units, which have a larger variety of responses than those of the MUAe (see Chapter 5). Future work will address these outstanding issues.

In conclusion, we found that many aspects of V1 MUA response to rapidly changing stimuli that cannot be understood in terms of a simple feed-forward process, can instead be partially explained as a consequence of recurrent connections within V1.

Chapter 5. Reverse connectivity simulations

5.1 Introduction

Many models of orientation selectivity in primary visual cortex assume a "Mexican Hat" type connectivity (see Figure 3b for an example) among neurons with different preferred orientations within a hypercolumn of visual cortex (neurons with overlapping receptive fields in visual space), as well as in center-surround interactions (Ben-Yishai et al., 1995; Somers et al., 1995; Carandini and Ringach, 1997; Teich and Qian, 2003). However, although some groups have been successful at obtaining rough physiological evidence of this aspect (Michalski et al., 1983; Ko et al., 2011), these type of in-vivo measurements are very hard to obtain even for one neuron, and thus even harder to generalize across neuronal populations. Specifically, neurons are likely to have idiosyncratic patterns of connectivity, which may explain some of the differences in their individual responses. The ability to measure the variability and patterns of intracortical connectivity would greatly enhance our mechanistic understanding of the cortical computations that give rise to perception of orientation. In the current chapter, we explore the possibility that we could combine modeling similar to that in Chapter 2 with electrophysiological measurements such as those in Chapter 4 to infer specific patterns of network connectivity for individual neurons.

In previous chapters we showed that for rapidly presented oriented gratings, stimulus history induces shifts in orientation tuning in primary visual cortex (Figure 7 and Figure 16). Through model simulations we confirmed that changing the connectivity profile leads to patterns of shifts with different magnitude and time courses. In addition, we found examples in our V1 recordings of units that have shifts of not only different magnitudes and different degree of symmetry (Figure 16, magnitudes not shown), but also simultaneously recorded single units, one with repulsive shifts and another with attractive shifts in their preferred orientation (Figure 24). This stands in difference to the findings in the multi-unit activity envelope, which always exhibited repulsive shifts (Figure 16). Note that our experimental data only served as motivation, and all data shown in this chapter apart from Figure 24 is model data.

We hypothesized that the pattern of shifts of any one V1 neuron is at least in part determined by its connectivity to other V1 neurons with different orientation preference. This led us to ask whether it's possible to use the neural response to stimulus sequences in combination with the model presented in Chapter 2 to infer the connectivity of a given neuron. Before attempting this with experimental data, we started with a proof of principle using model simulations as our "ground truth": taking the responses of a certain model to sequences of gratings as our "experimental data", can we reliably estimate the parameters of the connectivity profile in the underlying model (from a parameter space exploration)?

5.2 Methods

The Methods for the recordings from primary visual cortex shown in Figure 24 are described in Chapter 4. In this case we obtained action potentials of single isolated units using the automated clustering algorithm KlustaKwik (Harris et al., 2000).

For the connectivity estimates we first generated "experimental data" by simulating the response of the C-model described in Chapter 2 to sequences of gratings. We then tried to best match those responses while varying the four parameters that determine connectivity: cortical strength, ratio of inhibition to excitation and width of the excitatory and inhibitory lateral connectivity profiles. The search for the best model was an exhaustive search through a grid of the four connectivity parameters, with values (minimum:stepsize:maximum): $J_{cortex} = [0.9:0.1:1.8] \text{ mV/Hz}$; $r_{IE} = [1:0.02:1.18]$; $\kappa_E = [1.5:0.1:2.4]$; $\kappa_I = [1:0.1:1.9]$. The rest of the model parameters were set to the values of the C-model: $\tau = 10.8 \text{ ms}$; $\alpha = 10.6 \text{ Hz/mV}$; $J_{1gn} = 9.57 \text{ mV/Hz}$; $\kappa_{1gn} = 1.56$. Please note that the exact connectivity values of the C-model ($J_{cortex} = 1.71 \text{ mV/Hz}$; $r_{IE} = 1.18$; $\kappa_E = 1.59$; $\kappa_I = 1.16$) were only approximated by the grid, but not included in the exhaustive search.

To mimic real experimental conditions, each of these models, including the C-model, was presented with 1000 random sequences of ten orientations presented for 20 ms each, and their responses in time for each sequence were recorded at 1000 Hz (1 ms resolution). We then calculated the coefficient of determination, or proportion of explained variance for each model j as a measure of how well it fit the data:

$$R_j^2 = 1 - \frac{SS_{residuals}^j}{SS_{total}}, \text{ with } SS_{total} = \sum_i (c_i - \overline{c})^2 \text{ and } SS_{residuals}^j = \sum_i (m_i^j - c_i)^2$$
(9)

where m_i^j are the responses of model j, c_i are the responses for the C-model, and \overline{c} is the mean response of the C-model over time and over all sequences. If the sum of squares of the residuals is as large (or larger) than the total sum of squares for the C-model, then the model did not explain any more of the variance than the mean, and thus did not represent a good fit. Therefore we only considered models j for which $R_j^2 > 0$. This measure quantifies how well the model captured the dynamics of the simulated data. In a real experiment this would be the only measure of goodness-of-fit available to us. In the current simulations, however, we had access to the ground truth underlying connectivity, hence we also determined the same index (proportion of explained variance) to compare the connectivity profile of the C-model (the "experimental data") and the models.

5.3 Results

5.3.1 Motivation

We recorded from primary visual cortex using chronically implanted arrays (see Chapter 4 for details), and analyzed single unit responses. We found that even within one day of recording not all single units would have the same pattern of orientation tuning curve shifts. The magnitude was quite different for different neurons (not shown), the degree of symmetry varied (not shown), and most notably we found units for which the shifts in preferred orientation were attractive (Figure 24, red), rather than repulsive (Figure 24, blue) relative to the adapter orientation. We hypothesized that the pattern of shifts of single-units is (at least in part) determined by their profile of connectivity to other V1

neurons. Therefore we explored the possibility of inferring aspects of the connectivity profile by observing the effect of the recurrent interactions on the neural response.



Figure 24 - Examples of different patterns of shifts for simultaneously recorded single units in V1. **a:** Tuning curves of two simultaneously recorded units. The preferred orientation of the top (blue) curve was 6.5° and that of the bottom (red) curve 160° . **b:** Shifts in preferred orientation for the 12 adapting orientations. The top neuron (blue) exhibits repulsive shifts, whereas the bottom neuron (red) exhibits attractive shifts in preferred orientation. Error bars represent the standard deviation of the μ parameter of the von Mises fit for 1000 bootstrap samples.

5.3.2 Proof of principle

We performed simulations using the model presented in Chapter 2. Our aim was to advance towards a proof of principle: that finding a model that minimizes the difference in the response to some experimental data is a good method to infer aspects about neuronal properties that are not easy to measure physiologically. More precisely, we were interested in inferring aspects related to connectivity. Note that the experimental data shown in Figure 24 only served as motivation, but all analyses were performed on simulated data.



Figure 25 - The model that best matches the response also results in a close match to the connectivity profile, but is not necessarily the closest model to the "real" model in parameter space. **a:** Response of a neuron with preferred orientation 0° to a sequence of ten orientations shown for 20 ms each: -34, -17, 22, -16, -67, 17, 50, 98, -8, -7. The black curve corresponds to the C-model parameters, which we chose as our experimental data. Shown in green is the model that produces the smallest mean squared error of the responses. In blue is the model that takes the parameters in the grid closest to the C-model parameters. **b:** Connectivity profile of the three models in **a**. Note that the green and black profiles are nearly identical, but the blue is fairly different. This shows that distance in connectivity parameter space does not equate to difference in the connectivity profile.

We took the C-model presented in Chapter 2 to represent our experimental data, and performed an exhaustive parameter space exploration for the model parameters that determine the connectivity profile. We then presented a number of sequences of randomly selected orientations to both the C-model and all the models from the parameter space exploration, and saved the responses of each one. We then determined the model that resulted in the best fit to the C-model response by minimizing the squared error of the responses, and compare such model's properties with those of the C-model. The resulting model matched the response in time to a sequence of orientations very closely (99.93% of variance explained, compare black and green curves in Figure 25b). However, contrary to our intuition, the location of this model in the 4-dimensional connectivity parameter space was not the closest in distance to the real parameters of the C-model.

The similarity that we can obtain in both the responses and the underlying connectivity depend on how finely we are exhaustively exploring parameter space. As we showed in Figure 25, it is not possible to find one unique connected component in the parameter space grid that will assure similarity in the connectivity profile. Our next approach was to study whether it was necessarily true that a model that provides a good fit to the responses to the grating sequences, also provides a good fit in terms of the connectivity profile, regardless of its location in the parameter space grid. To tackle this question we calculated the coefficient of determination to quantify the proportion of variance of the experimental data that each model explained (see Methods), and the proportion of

variance of the ground truth connectivity profile explained by the connectivity profiles of each of the models. We only considered those models that explained a positive fraction of the variance of the C-model time course and plotted the proportion explained variance as a function of the fraction explained variance of the connectivity profile (Figure 26). We found that explaining a positive fraction of the variance of the time course does not necessarily imply a good fit to the connectivity profile, as illustrated by all the models that have a negative index of explained variance for the connectivity.



Figure 26 - Models that match the response of the experimental data don't necessarily have similar connectivity profiles. Each blue dot corresponds to a model resulting from one point in the parameter space exploration, such that it explained a positive fraction of the variance of the time course data (see **Methods** for details). The explained variance for the connectivity profiles yields an index for how well the profiles of the model match the ground truth profile. An R^2 of 1 represents a perfect match, and any value below zero represents a model that is worse than a model of constant (all to all) connectivity.

5.3.3 Robustness to noise

Given that in reality experimental data will have measurement noise, we also tested the robustness of this approach. We did this by injecting Poisson-like multiplicative Gaussian noise to the C-model (see Methods) and repeating the proportion of variance explained analysis. We found that the best fitting model was exactly the same as the one found above, the difference was that the proportion of variance explained of the time course went down to 88.67% (Figure 27).



Figure 27 - Adding noise to the C-model does not impair the ability to find the model with the most similar connectivity profile. **a:** The model that produces the smallest error in the time course when compared to the C-model with added noise is the model that has the most similar connectivity profile, the same shown in **Figure 25**. **b:** The proportions of variance explained of the C-model with added noise are similar to those of the C-model with no noise, except lower numbers are attained for the Response \mathbb{R}^2 .

5.4 Discussion

We found that fitting our model to the response to several sequences of randomly presented orientations using a least squares approach results in a good approximation of the response time course (99.93% variability explained), even when the data were noisy

(88.67% variability explained, Figure 25 and Figure 27). The main finding of this computational analysis, however, was that a model that explains a large proportion of the variance of the responses may not have a similar connectivity profile. For example there are models that explain around 80% of the variance of the time course (with no noise), but fare very poorly explaining the variance of the connectivity profile (Figure 26).

This implies that measuring the responses to sequences of gratings with the same contrast, and the same duration for each grating orientation, does not provide a strong enough constraint to the model. In other words, it is too easy for a model to have a very similar time course under the tested conditions, without the connectivity being similar. Several examples of similar over-representation of parameter space are reviewed in Marder and Taylor (2011).

As an example, one of the models that explained 80% of the variance of the time course had a net inhibitory connectivity for all orientations, which caused the response to fall off shortly after a transient rising. Because our gratings were only presented for 20 ms each, this did not prevent the model from fitting the time course of our C-model well, but it would have explained a lot less of the variance had we presented some gratings for longer. A natural next step would be then to use sequences of gratings with varying durations and/or contrasts to constrain the model further, as well as add blank periods to constrain the fall of the response.

A limitation of this study is that the connectivity profile of neurons in primary visual cortex is probably not limited to Mexican hats. In Figure 24 we showed that some

neurons have attractive, rather than repulsive shifts, which could not be obtained with the traditional Mexican hat connectivity. Therefore our analysis should be expanded to include other forms of lateral connectivity that would lead to these patterns of shifts, and confirm that in those cases our approach also leads to good estimates of connectivity.

Measuring effective connectivity in-vivo is a long standing problem in Neuroscience. Recent advances have been made at a gross level with the help of imaging techniques such as functional imaging combined with computational approaches such as Granger causality and dynamic causal modeling (Friston, 2011); that use correlation between simultaneously recorded units (or voxels) to determine connectivity. This is very different to the model based approach described here, but with the advent of chronically implanted multi-electrode arrays, the possibility of using both approaches at the same time to gain power is very exciting.

In conclusion, even though this preliminary finding is a negative finding as of now, it has been very important to help us find out, for example, whether it is feasible to collect enough data in one recording session to attempt a reverse connectivity inference. We also learned that we need to add more variability in the duration or contrast of the gratings to constrain the models better. This is something we would not have known without the model, and if we'd just gone ahead and tried doing the experiment, it would not have been valid.

Chapter 6. General discussion

"Adaptation", in its multiple forms and meanings, has mostly been studied at time-scales of several seconds, to minutes, to even hours and days (in particular with behavioral and modeling approaches). But humans, as many other mammals, make several eye movements per second during normal exploratory behavior (Carpenter, 1988), so it is actually the timescale of a few dozen to a couple hundred milliseconds that is most relevant for behavior. Furthermore, it has been suggested that visual processing of a complex natural image can be achieved in less than 150 ms (Thorpe et al., 1996); consistent with the finding that most of the information about a constant visual stimulus (even constant motion) is carried in the first 250 ms of the sensory response (Buracas et al., 1998).

We have shown in Chapter 2 that recurrent connections within a realistic network model of primary visual cortex can generate reverberations in neural activity for up to 1 second (Figure 8). This suggests that recurrence, a prominent characteristic of cortical networks, is acting precisely in the most ecologically relevant time window. This finding highlights the logical flaw of concluding that any "slow" change in the neural response or selectivity lasting more than the membrane time constant, is necessarily due to plasticity. On the contrary, dynamic neural systems similar to those modeled here are almost certain to exhibit changes in tuning at these timescales as responses evolve through the "pushing and pulling" of recurrent interactions. The persistence of such effects would depend on the particular configuration of connectivity; in theory, recurrent networks can maintain patterns of activity indefinitely without further input if appropriately parameterized (Wang, 2001).

One of the most exciting aspects of the recurrent model presented in Chapter 2 is its ability to make testable predictions, and the ease of subjecting the model to different experiments and analyses now that the base code is implemented (for example, Chapter 5). It is important to keep in mind, however, that often our intuition about how even a simple network behaves can be wrong or misleading. Therefore we advise to always execute simulations before making inferences about how the network would behave in different situations. As noted above, previous investigators have assumed that the standard recurrent model would not generate the shifts they observed and proceeded to add plasticity parameters, such as a re-weighting of the connectivity profile. The model that Teich and Qian (2003) used, for example, is almost identical to that used here and our implementation of their model (data not shown) showed shifts similar to those reported here without any plasticity.

Another attractive perspective of the model is how it helps us to design better experiments by testing them on the model before collecting the data. As an example, our analysis in Chapter 5 showed that presenting sequences of orientations with the same duration and contrast may not be enough to capture the intricacies of the recurrent connectivity in the responses. In the same way the duration of the test stimulus in Chapter 3 was chosen based on our knowledge of the time course of the shifts in tuning curves obtained from model simulations fit to neural data (Figure 7c). The main limitation in this work is that our hypothesis that the shifts in orientation tuning curves as a consequence of rapidly presented oriented stimuli are due to population dynamics, rather than plasticity, is currently not testable. To our knowledge, as of now there are no in-vivo methods to switch off or prevent plasticity and test if these shifts still occur. Therefore we do not (and cannot) claim that there is no plasticity contributing to this effect, but we do emphasize that recurrent connections are at least part of the story, if not the whole story.

Another contributing factor to the shifts in orientation tuning at fast timescales could be differences in latencies among the inputs to V1, for example parvocellular vs. magnocellular LGN neurons. If the latency in the V1 response was determined by the magnocellular input, and the parvocellular input was delayed by, say, x ms; then for the first x ms of what we're considering the response to the test orientation, the parvocellular input would correspond to the adapter orientation. However, the distribution of magnocellular and parvocellular latencies have been found to be mostly overlapping, and the difference between the earliest responses of both types of neurons was found to be less than 10 ms (Maunsell et al., 1999). Recordings from V1 show that even excluding the first 100 ms from the response to the test, the shift in orientation tuning curves is still present (see Figure 7d and section 2.2.3) (Patterson et al., 2013). Therefore different input latencies could contribute, but not fully account for, the tuning shifts discussed in this document.

Furthermore, different axonal conduction speeds can only explain about half of the delay of parvocellular neurons in reaching cortex (Dreher et al., 1976; Schiller and Malpeli,

1977, 1978; Kaplan and Shapley, 1982). To account for the remaining delay, other factors have been proposed, such as differences in retinal circuits feeding each channel, as well as the effectiveness of stimulus intensities in driving each channel; but these have not been explained fully (Maunsell et al., 1999). Another possibility is a difference in the dynamics of parvocellular vs. magnocellular channels. For example, in the model presented in Chapter 2, longer latencies arise when the lateral connections are stronger (Figure 9). The remaining delay could therefore very well be due to stronger recurrent connectivity in the pathway leading to parvocellular neurons with respect to magnocellular neurons, although this is currently speculation.

Recurrent connectivity as an underlying explanation for the shifts in orientation tuning could also help solve some outstanding controversies. As an example, Dragoi et al. (2001) found that shifts in orientation preference are largest in magnitude near pinwheel centers, concluding that plasticity is strongest in those locations of visual cortex. In the same year, Schuett et al. (2001) showed that plasticity in orientation tuning can be induced by pairing a brief oriented stimulus with electrical stimulation in the cortex. Further, the directions of the resulting shifts in orientation tuning (attractive or repulsive) were determined by the relative timing of visual and electrical stimulation. The interesting part for us is that they found that this pairing-induced plasticity is minimal in pinwheel centers versus other areas of the orientation map, in apparent contradiction to the Dragoi et al. (2001) result. One possible explanation could be that adaptation through visual stimulation generates a different kind of plasticity than that generated by the paired visual and electrical stimulation. But a more parsimonious explanation is that the shifts

observed by Dragoi et al. (2001) are perhaps mostly due to reverberations in the recurrent network, as we have shown in Chapter 2.

In Chapter 2 we showed that attractive shifts in the population response (and therefore in perception, under a labelled line assumption) (Figure 5) generate repulsive shifts in tuning curves (Figure 6). Our behavioral experiments (Chapter 3) confirmed the attractive nature of the perceptual shifts on a short time scale. The consistency between model predictions and behavioral outcomes supports our view that network dynamics are critical for understanding neural information processing.

In principle the behavioral experiments shown in Chapter 3 and the physiological recordings shown in Chapter 4 could be done concurrently on macaques, to be able to correlate the perception of the animal with neural activity in V1 directly. However, given our efforts to minimize plasticity, our behavioral task is somewhat counter-intuitive, needing lots of practice and explaining, that would make it difficult to train a monkey to perform well. In addition, lots of repetitions are needed to accurately quantify neural responses, which would take too long to obtain if we only get one adaptation effect per trial. In this sense presenting gratings in rapid streams is ideal for obtaining multiple repetitions of all transitions within a recording session, but cannot be easily adapted to a behavioral task. Therefore we performed our neural recordings on monkeys and used human participants for the behavioral studies. This is a clear example of how human and animal studies are both important, and can complement each other to optimize resources.

In addition, it is important to keep in mind differences across species (human, macaque, cat, rodents), as well as across awake and anesthetized states, when combining studies to make general inferences. For example, rodents have a salt-and-pepper distribution of preferred orientation in primary visual cortex, unlike the smooth orientation maps that have been found in cats, primates and humans (Kaschube, 2014). However, the intracortical connectivity for primary visual cortex neurons seems to be similar with both types of organization: mostly isotropic, but slightly more frequent among neurons with similar preferred orientations (Michalski et al., 1983; Ko et al., 2011). This raises the interesting possibility of using our reverse connectivity approach to study functional connectivity in these two very different types of organization.

The effects of recurrent reverberations on neural tuning and computations are not limited to visual processing in V1, nor even to sensory processing more broadly. Recurrent connectivity and persistence are ubiquitous in the brain and are likely to be crucial for almost all aspects of brain function. Further, the timescales of recurrent persistence are not limited to tens to hundreds of milliseconds, but rather can be on the order of seconds or even minutes in some cases. A clear example is in the domain of working memory for cognition and action planning. In parietal cortex, for example, activity related to planned eye and arm movements can be maintained for several seconds (at least) after a target stimulus is removed (Gnadt and Andersen, 1988; Snyder et al., 1997; for a review, see Andersen and Buneo, 2002). This sustained activity is likely to underlie the ability to remember a location of interest across a delay period for the purpose of goal-directed movement and would have to be achieved through recurrent feedback loops. This

memory-related activity serves a clear purpose, but analogous to the tuning shifts in our V1 model, it would also bias how the population responds to new incoming information and cause shifts in the spatial tuning of individual neurons. Similar history effects would also occur at faster timescales in the domain of action-planning and would have to be taken into account to have a full understanding of the cortical representations for action. This would be important, for example, for the purpose of designing accurate neural prostheses for the control of robotic limbs.

6.1 Conclusions

The work described in this thesis used three very distinct but highly complementary approaches to study the dynamics of the representation of orientation in primary visual cortex: computational modeling, human psychophysics, and alert animal electrophysiology. Using computational modeling we were able to identify previously unknown properties of standard recurrent models of orientation representation. We showed that these properties – adaptation-like shifts in tuning curves – matched very closely electrophysiological observations from the literature and offered a more parsimonious explanation for this behavior than previous accounts. We then translated our observations from the modeling domain into the realm of human psychophysics to draw links between neural representations and human perception. There, we showed that the perception of orientation in dynamic sensory environments matched that predicted from the dynamics of population activity in our model of V1. We then probed neural activity in V1 directly using electrophysiological methods in alert, behaving primates to further examine the influence of recurrent connectivity on the representation of orientation. We showed that multi-unit activity in V1 responds non-linearly to dynamic orientation stimuli and that many of the key observations were captured by the recurrent orientation model. Finally, we offered preliminary steps toward a method to infer the connectivity of experimentally observed neurons by combining dynamic orientation stimuli with model-fitting.

We believe that it is the combination of computational, behavioral, and/or electrophysiological techniques that will lead to the most interesting breakthroughs. None of these techniques have to stand on their own, as each can only provide circumstantial evidence, but it is the combination and consistency of the findings that makes a strong case.

References

- Addams R (1834) An account of a peculiar optical phenomenon seen after having looked at a moving body, etc. London & Edinburgh Philosophical Magazine and Journal of Science 5:373-374.
- Andersen RA, Buneo CA (2002) Intentional Maps in Posterior Parietal Cortex. Annu Rev Neurosci 25:189-220.
- Asari H, Zador AM (2009) Long-lasting context dependence constrains neural encoding models in rodent auditory cortex. J Neurophysiol 102:2638-2656.
- Ben-Yishai R, Bar-Or RL, Sompolinsky H (1995) Theory of orientation tuning in visual cortex. Proc Natl Acad Sci USA 92:3844-3848.
- Benucci A, Ringach DL, Carandini M (2009) Coding of stimulus sequences by population responses in visual cortex. Nature Neurosci 12:1317-1324.
- Buracas GT, Zador AM, DeWeese MR, Albright TD (1998) Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. Neuron 20:959-969.
- Burnham KP, Anderson DR (2002) Model selection and multi-model inference: A practical Information-theoretic approach, 2nd Edition. New York: Springer-Verlag.
- Burr DC, Wijesundra SA (1991) Orientation discrimination depends on spatial frequency. Vision Res 31:1449-1452.
- Campbell FW, Maffei L (1970) Electrophysiological evidence for the existence of orientation and size detectors in the human visual system. The Journal of physiology 207:635-652.
- Campbell FW, Maffei L (1971) The tilt after-effect: a fresh look. Vision Res 11:833-840.
- Carandini M, Ringach DL (1997) Predictions of a recurrent model of orientation selectivity. Vision Res 37:3061-3071.
- Carandini M, Ferster D (2000) Membrane potential and firing rate in cat primary visual cortex. J Neurosci 20:470-484.
- Carpenter RHS (1988) Movements of the eyes, 2nd Edition. London: Pion.

- Chung S, Ferster D (1998) Strength and orientation tuning of the thalamic input to simple cells revealed by electrically evoked cortical suppression. Neuron 20:1177-1189.
- Clifford CW, Wyatt AM, Arnold DH, Smith ST, Wenderoth P (2001) Orthogonal adaptation improves orientation discrimination. Vision Res 41:151-159.
- Compte A, Brunel N, Goldman-Rakic PS, Wang XJ (2000) Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. Cereb Cortex 10:910-923.
- DeValois RL, Yung EW, Hepler N (1982) The orientation and direction selectivity of cells in macaque visual cortex. Vision Res 22:531-544.
- Dragoi V, Sharma J, Sur M (2000) Adaptation-induced plasticity of orientation tuning in adult visual cortex. Neuron 28:287-298.
- Dragoi V, Rivadulla C, Sur M (2001) Foci of orientation plasticity in visual cortex. Nature 411:80-86.
- Dragoi V, Sharma J, Miller EK, Sur M (2002) Dynamics of neuronal sensitivity in visual cortex and local feature discrimination. Nature Neurosci 5:883-891.
- Dreher B, Fukada Y, Rodieck RW (1976) Identification, classification and anatomical segregation of cells with X-like and Y-like properties in the lateral geniculate nucleus of old-world primates. J Physiol 258:433-452.
- Duong T, Moore BDt, Freeman RD (2011) Adaptation changes stereoscopic depth selectivity in visual cortex. The Journal of neuroscience : the official journal of the Society for Neuroscience 31:12198-12207.
- Felsen G, Shen YS, Yao H, Spor G, Li C, Dan Y (2002) Dynamic modification of cortical orientation tuning mediated by recurrent connections. Neuron 36:945-954.
- Ferster D, Miller KD (2000) Neural Mechanisms of Orientation Selectivity in the Visual Cortex. Annu Rev Neurosci 23:441-471.
- Fischer J, Whitney D (2014) Serial dependence in visual perception. Nature Neurosci 17:738-743.
- Friston KJ (2011) Functional and effective connectivity: a review. Brain connectivity 1:13-36.
- Gardner JL, Anzai A, Ohzawa I, Freeman RD (1999) Linear and nonlinear contributions to orientation tuning of simple cells in the cat's striate cortex. Visual Neurosci 16:1115-1121.
- Gibson JJ, Radner M (1937) Adaptation, aftereffect and contrast in the perception of tilted lines: I. Quantitative studies. J Exp Psychol 20:453-467.
- Gnadt JW, Andersen RA (1988) Memory related motor planning activity in posterior parietal cortex of macaque. Exp Brain Res 70:216-220.
- Harris JP, Calvert JE (1989) Contrast, spatial frequency and test duration effects on the tilt aftereffect: implications for underlying mechanisms. Vision Res 29:129-135.
- Harris KD, Henze DA, Csicsvari J, Hirase H, Buzsaki G (2000) Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. J Neurophysiol 84:401-414.
- He S, MacLeod DI (2001) Orientation-selective adaptation and tilt after-effect from invisible patterns. Nature 411:473-476.
- Hegde J (2009) How Reliable is the Pattern Adaptation Technique? a Modeling Study. J Neurophysiol.
- Hubel DH, Wiesel TN (1962) Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. J Physiol 160:106-154.
- Kaplan E, Shapley RM (1982) X and Y cells in the lateral geniculate nucleus of macaque monkeys. J Physiol 330:125-143.
- Kaschube M (2014) Neural maps versus salt-and-pepper organization in visual cortex. Curr Opin Neurobiol 24:95-102.
- Ko H, Hofer SB, Pichler B, Buchanan KA, Sjostrom PJ, Mrsic-Flogel TD (2011) Functional specificity of local synaptic connections in neocortical networks. Nature 473:87-91.
- Kohn A (2007) Visual Adaptation: Physiology, Mechanisms, and Functional Benefits. J Neurophysiol 97:3155-3164.
- Kontsevich LL, Tyler CW (1999) Bayesian adaptive estimation of psychometric slope and threshold. Vision Res 39:2729-2737.
- Krekelberg B, van Wezel RJ, Albright TD (2006a) Adaptation in macaque MT reduces perceived speed and improves speed discrimination. J Neurophysiol 95:255-270.
- Krekelberg B, Boynton G, Vanwezel R (2006b) Adaptation: from single cells to BOLD signals. Trends Neurosci 29:250-256.
- Kristjansson A (2011) The functional benefits of tilt adaptation. Seeing Perceiving 24:37-51.

- Lamme VA, Roelfsema PR (2000) The distinct modes of vision offered by feedforward and recurrent processing. Trends Neurosci 23:571-579.
- Mach E (1861) Über das Sehen von Lagen and Winkeln durch die Bewegung des Auges. Sitzungsberichte der Kaiserlichen Akademie der Wissenschaften 43:215-224.
- Maffei L, Campbell FW (1970) Neurophysiological localization of the vertical and horizontal visual coordinates in man. Science 167:386-387.
- Magnussen S, Johnsen T (1986) Temporal aspects of spatial adaptation. A study of the tilt aftereffect. Vision Res 26:661-672.
- Marder E, Taylor AL (2011) Multiple models to capture the variability in biological neurons and networks. Nat Neurosci 14:133-138.
- Maunsell JH, Ghose GM, Assad JA, McAdams CJ, Boudreau CE, Noerager BD (1999) Visual response latencies of magnocellular and parvocellular LGN neurons in macaque monkeys. Vis Neurosci 16:1-14.
- Michalski A, Gerstein GL, Czarkowska J, Tarnecki R (1983) Interactions between cat striate cortex neurons. Exp Brain Res 51:97-107.
- Mitchell DE, Muir DW (1976) Does the tilt after-effect occur in the oblique meridian? Vision Res 16:609-613.
- Morant RB, Harris JR (1965) Two Different after-Effects of Exposure to Visual Tilts. The American journal of psychology 78:218-226.
- Muir D, Over R (1970) Tilt aftereffects in central and peripheral vision. J Exp Psychol 85:165-170.
- Muller JR, Metha AB, Krauskopf J, Lennie P (1999) Rapid adaptation in visual cortex to the structure of images. Science 285:1405-1408.
- Nelson S, Toth L, Sheth B, Sur M (1994) Orientation selectivity of cortical neurons during intracellular blockade of inhibition. Science 265:774-777.
- Paradiso MA, Meshi D, Pisarcik J, Levine S (2012) Eye movements reset visual perception. J Vis 12:11.
- Patterson CA, Wissig SC, Kohn A (2013) Distinct effects of brief and prolonged adaptation on orientation tuning in primary visual cortex. J Neurosci 33:532-543.
- Priebe NJ, Ferster D (2008) Inhibition, spike threshold, and stimulus selectivity in primary visual cortex. Neuron 57:482-497.

Rajimehr R (2004) Unconscious orientation processing. Neuron 41:663-673.

- Sanchez-Vives MV, Nowak LG, McCormick DA (2000) Membrane mechanisms underlying contrast adaptation in cat area 17 in vivo. J Neurosci 20:4267-4285.
- Schiller PH, Malpeli JG (1977) Properties and tectal projections of monkey retinal ganglion cells. J Neurophysiol 40:428-445.
- Schiller PH, Malpeli JG (1978) Functional specificity of lateral geniculate nucleus laminae of the rhesus monkey. J Neurophysiol 41:788-797.
- Schuett S, Bonhoeffer T, Hubener M (2001) Pairing-induced changes of orientation maps in cat visual cortex. Neuron 32:325-337.
- Schummers J, Cronin B, Wimmer K, Stimberg M, Martin R, Obermayer K, Koerding K, Sur M (2007) Dynamics of orientation tuning in cat v1 neurons depend on location within layers and orientation maps. Front Neurosci 1:145-159.
- Schwartz O, Hsu A, Dayan P (2007) Space and time in visual context. Nature Rev Neurosci 8:522-535.
- Sclar G, Lennie P, DePriest DD (1989) Contrast adaptation in striate cortex of macaque. Vision Res 29:747-755.
- Sillito AM, Kemp JA, Milson JA, Berardi N (1980) A re-evaluation of the mechanisms underlying simple cell orientation selectivity. Brain Res 194:517-520.
- Snyder LH, Batista AP, Andersen RA (1997) Coding of intention in the posterior parietal cortex. Nature 386:167-170.
- Solomon SG, Peirce JW, Dhruv NT, Lennie P (2004) Profound contrast adaptation early in the visual pathway. Neuron 42:155-162.
- Somers DC, Nelson SB, Sur M (1995) An emergent model of orientation selectivity in cat visual cortical simple cells. J Neurosci 15:5448-5465.
- Sompolinsky H, Shapley R (1997) New perspectives on the mechanisms for orientation selectivity. Curr Opin Neurobiol 7:514-522.
- Super H, Roelfsema PR (2005) Chronic multiunit recordings in behaving animals: advantages and limitations. Prog Brain Res 147:263-282.
- Teich AF, Qian N (2003) Learning and adaptation in a recurrent model of V1 orientation selectivity. J Neurophysiol 89:2086-2100.

- Thorpe S, Fize D, Marlot C (1996) Speed of processing in the human visual system. Nature 381:520-522.
- Ulanovsky N, Las L, Farkas D, Nelken I (2004) Multiple time scales of adaptation in auditory cortex neurons. The Journal of neuroscience : the official journal of the Society for Neuroscience 24:10440-10453.
- Wang XJ (2001) Synaptic reverberation underlying mnemonic persistent activity. Trends Neurosci 24:455-463.
- Wenderoth P, van der Zwan R (1989) The effects of exposure duration and surrounding frames on direct and indirect tilt aftereffects and illusions. Percept Psychophys 46:338-344.
- Wichmann FA, Hill NJ (2001a) The psychometric function: II. Bootstrap-based confidence intervals and sampling. Percept Psychophys 63:1314-1329.
- Wichmann FA, Hill NJ (2001b) The psychometric function: I. Fitting, sampling, and goodness of fit. Percept Psychophys 63:1293-1313.
- Wissig SC, Kohn A (2012) The influence of surround suppression on adaptation effects in primary visual cortex. J Neurophysiol 107:3370-3384.
- Wolfe JM (1984) Short test flashes produce large tilt aftereffects. Vision Res 24:1959-1964.
- Zhao C, Series P, Hancock PJ, Bednar JA (2011) Similar neural adaptation mechanisms underlying face gender and tilt aftereffects. Vision Res 51:2021-2030.